

ip interface detail Displays detailed information about the specified interfaces or all interfaces.

✓ 3500
9000
9400

Valid Minimum Abbreviation

`ip i det`

Important Consideration

- When you enter the command, you are prompted for an interface index number even if you have only one interface defined.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the IP interface whose summary information you want to display	<ul style="list-style-type: none"> ■ One or more configured indexes ■ all ■ ? (for a list of selectable indexes) 	–

Fields in the IP Interface Detail Display

Field	Description
ARP proxy	Whether ARP proxy is enabled or disabled for the specified interface.
Broadcast address	Broadcast address for the specified interface.
Directed broadcast	Whether the forwarding of a directed broadcast (all 1s in the host portion of the address) is enabled or disabled for the specified interface. (A directed broadcast is a packet that is sent to a specific network or series of networks.)
ICMP redirect	Whether ICMP redirect is enabled or disabled for the specified interface.
ICMP router discovery	Whether the ICMP Router Discovery is enabled or disabled for the specified interface
Index	Index number that is associated with the interface.
IP address	IP address of the interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.

Field	Description
Preference	Whether there is a preference being used for the specified interface. If ICMP router discovery is enabled, the system uses the routing interface with the highest preference level.
State	State of the IP interface. It indicates whether the interface is available for communications (<code>up</code>) or unavailable (<code>down</code>).
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnetwork number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.
Type	Type of interface: VLAN-based (<code>VLAN</code>) or router port-based (<code>port</code>).
Index	Index number of the IP VLAN that is associated with the IP interface.
MaxAdvInterval	Maximum advertisement interval between ICMP router discovery advertisements (in seconds).
MinAdvInterval	Minimum advertisement interval between ICMP router discovery advertisements (in seconds).
Holdtime	Length of time that ICMP router discovery advertisements are held valid.
State	State of the IP interface. It indicates whether the interface is available for communications (<code>up</code>) or unavailable (<code>down</code>).
ID	<ul style="list-style-type: none"> ■ Logical port number of the IP interface (if the Type field displays <code>port</code>) ■ VLAN index number that is associated with the IP interface (if the Type field displays <code>VLAN</code>)

IP Interface Detail Example (3500)

```
Select menu option (ip/interface): detail
Select IP interfaces (1|all|?) [1]: 1
IP routing is disabled
```

Index	IP address	Subnet mask	State	Type	ID
1	158.101.31.21	255.255.255.0	Down	Port	1

Index	ARP proxy	Broadcast address	Directed broadcast	ICMP redirect
1	enabled	255.255.255.255	enabled	enabled

Index	ICMP router discovery	Preference	MaxAdvInterval	MinAdvInterval	Holdtime
1	disabled	n/a	n/a	n/a	n/a

**ip interface define
(3500/9000 Layer 3)****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Defines an IP interface.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**`ip i def`**Important Considerations**

- When you define an IP interface, you must decide whether you want the interface to use router port-based routing or VLAN-based routing.
 - Router port-based routing directs the system to attempt to route the frame before it attempts to bridge the frame.

When you set up a router port-based IP interface, the system automatically creates a virtual LAN (VLAN) for the interface. The system assigns the next available VLAN index number to this VLAN.

- VLAN-based routing directs the system to attempt to bridge the frame before it attempts to route the frame.

When you set up a VLAN-based IP interface, you must first define a VLAN and select IP as a protocol supported by the VLAN, as described in Chapter 14.



If you define a router port, you do not have to define the VLAN first; the corresponding single-port VLAN is automatically defined.

- Port-based routing uses allClosed mode; VLAN-based routing uses either allClosed or allOpen mode. If you attempt to set up a router port-based IP interface in allOpen mode, the system notifies you with a message that it will change the VLAN mode to allClosed and recreate the default VLAN, clearing your existing VLANs in the process. Then the system prompts you to continue. (See the port-based router example at the end of this command description.)
- You cannot define a port-based IP interface on a port that is already a member of a VLAN-based IP interface. To change from one type of interface to another, you must redefine all IP interfaces and VLANs that are associated with that port.



CAUTION: *Using different routing models (port-based or VLAN-based) in the same network without careful planning can adversely affect your network operations. Be sure that you understand the potential effects of router port-based and VLAN-based routing on your network. See the Implementation Guide for the CoreBuilder 3500 and for the CoreBuilder 9000 for detailed information about IP interfaces and VLANs.*

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address of the interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.	A valid IP address in the range of addresses that are assigned to your organization	–
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnetwork number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.	A valid subnet mask in accordance with the bits that are used for network number, subnetwork, and host number	Depends on specified IP address
Interface type	Whether to use router port-based routing or VLAN-based routing.	<ul style="list-style-type: none"> ■ port ■ vlan 	vlan
VLAN mode (for router port-based routing)	Whether the system removes all VLANs and recreates the default VLAN to enable port-based routing.	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Bridge port (for router port-based routing)	Port to use for port-based routing (may designate only one port).	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of selectable ports) 	–
VLAN interface index (for VLAN-based routing)	Index number of the IP VLAN that is associated with the IP interface; for a VLAN-based IP interface, you must assign this number. (Not applicable if you have more than one VLAN)	<ul style="list-style-type: none"> ■ A selectable VLAN index ■ ? (for a list of selectable VLAN indexes) 	Next available index number

IP Interface Define Example (Port-based Routing)

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter interface type (vlan/port) [vlan]: port
VLAN mode must be changed to allClosed to support this
interface.
This removes all VLANs, then re-creates the Default VLAN.
continue? (n,y) [y]: y
Select bridge port (1-6|?): 1
```

IP Interface Define Example (VLAN-based Routing)

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter interface type (vlan/port) [vlan]:vlan
Enter VLAN interface index {3|?} [3]: 3
```

ip interface define
(3900/9300/9400/
9000 Layer 2)

Defines an IP interface.

Valid Minimum Abbreviation

`ip i def`

Important Consideration

- Before you define the IP (routing) interface, first define a virtual LAN (VLAN) and select IP as a protocol that the VLAN supports, as described in Chapter 16.

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address of the interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.	A valid IP address in the range of addresses that are assigned to your organization	–
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnetwork number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.	A valid subnet mask in accordance with the bits that are used for network number, subnetwork, and host number	Depends on specified IP address
VLAN interface index	Index number of the IP VLAN that is associated with the IP interface. (Not applicable if you have more than one VLAN)	<ul style="list-style-type: none"> ■ A selectable VLAN index ■ ? (for a list of selectable VLAN indexes) 	Current value

IP Interface Define Example

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter VLAN interface index {2|?} [2]: 2
```

- 3500
- ✓ 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

ip interface modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- ✓ 9400

Changes the configuration of an interface that you have already defined.

Valid Minimum Abbreviation

`ip i m`

- ✓ 3900
- ✓ 9300

Important Consideration

- On the CoreBuilder 3500, you cannot modify the port number (router port-based routing) after it has been defined because of the associated virtual LAN (VLAN); you must remove the interface and then redefine it.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number that is associated with the interface that you want to modify. (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ A selectable IP interface index ■ ? (for a list of selectable indexes) 	Current value
IP address	IP address of the interface that you want to modify.	A valid IP address in the range of addresses that are assigned to your organization	Current IP address
Subnet mask	Subnet mask for the interface that you want to modify.	A valid subnet mask in accordance with the bits that are used for network number, subnetwork, and host number	Current subnet mask
VLAN interface index (for VLAN-based routing)	Index number of the IP VLAN that is associated with the IP interface; for a VLAN-based IP interface, you must assign this number. (Not applicable if you have more than one VLAN)	<ul style="list-style-type: none"> ■ A selectable VLAN index ■ ? (for a list of selectable VLAN indexes) 	Current value

ip interface remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Removes an IP interface from the system's routing table.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

ip i re

✓ 3900

✓ 9300

Important Considerations

- Before you remove the interface, remove any static entries in the routing table or the Address Resolution Protocol (ARP) cache.
- On the CoreBuilder 3500, if you remove a router port-based IP interface, the system removes the virtual LAN (VLAN) that is associated with it as well.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number that is associated with the interfaces that you want to remove (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Current value

ip interface arpProxy *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

3900
9300

On a per-interface basis, enables or disables ARP proxy, which helps end stations on a subnetwork reach remote subnetworks that do not have routing capabilities or a default gateway configured.

Valid Minimum Abbreviation

ip i a

Important Considerations

- When ARP proxy is enabled and an end station sends an Address Resolution Protocol (ARP) request for a remote network, the system determines if it has the best route and then answers the ARP request by sending its own MAC address to the end station. The end station then sends the frames for the remote destination to the system, which uses its own routing table to reach the destination on the other network.
- When an interface is defined, the default ARP proxy state is `enabled`.
- The end stations must view the entire network configuration as one network (that is, by using a smaller subnet mask).
- Evaluate prolonged use of ARP proxy because it has some drawbacks, including increased ARP traffic and a need for larger ARP tables to handle the mapping of IP addresses to MAC addresses.

Options

Prompt	Description	Possible Values	[Default]
Interface	Index number for the interface for which you want to enable or disable ARP proxy. (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Current value
ARP proxy state	Whether you want to implement ARP proxy on an interface. The system prompts you for a state for each interface.	<ul style="list-style-type: none"> ■ enabled ■ disabled 	Current value

IP Interface ARP Proxy Example (3500)

```
Select menu option (ip/interface): arpproxy  
Select IP interfaces (1,2|?|all):2  
  Interface 2 - Enter proxy state (disabled, enabled)  
[enabled]: enabled
```

**ip interface
broadcastAddress**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

On a per-interface basis, assigns the broadcast address that the system uses to forward the received directed broadcasts and advertise Routing Information Protocol (RIP) packets.

Valid Minimum Abbreviation

ip i b

Important Considerations

- You assign the broadcast address on a per-interface basis.
- When an IP interface is configured, its default broadcast address is 255.255.255.255.
- The broadcast address that you specify affects the RIP advertisement address that is used for the RIP interface. You see the specified broadcast address as the advertisement address under the RIP menus. See "ip rip display" later in this chapter for information about the RIP interface display.
- You cannot change the broadcast address for an interface if you have added any RIP advertisement addresses to that interface. See "ip rip addAdvertisement" later in this chapter for more information.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces to which you want to assign a broadcast address (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ One or more interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Current value
Broadcast address per interface	Broadcast address that you want to assign to an interface	A valid address	Current address

ip interface directedBroadcast

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Specifies whether the forwarding of a directed broadcast (all 1s in the host portion of the address) is enabled or disabled for a specified interface. A *directed broadcast* is a packet that is sent to a specific network or series of networks.

Valid Minimum Abbreviation

ip i di

Important Considerations

- You define the directed broadcast state on a per-interface basis.
- When the state is enabled and the system determines that the destination is different from the interface that is receiving the directed broadcast, the system uses the broadcast address that is defined for this interface to forward the directed broadcast.
- You can disable the forwarding of a directed broadcast if security is an issue.
- By default, the directed broadcast state is enabled.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index numbers of the interfaces to which you want to enable or disable the forwarding of a directed broadcast. (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Current value
Directed broadcast state	Whether you want to implement direct broadcast on an interface. The system prompts you for a state for each interface.	<ul style="list-style-type: none"> ■ enabled ■ disabled 	Current value

IP Interface Directed Broadcast Example (3500)

```
Select menu option (ip/interface): directedBroadcast
Select IP interfaces (1,2|all|?):2
Interface 2 - Enter directed broadcast state
(disabled, enabled) [enabled]:
```

**ip interface
icmpRedirect**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables the transmission of an Internet Control Message Protocol (ICMP) redirect to the sender of a frame to indicate that there is a better gateway available to handle the frame than this routing interface.

Valid Minimum Abbreviation

```
ip i icmpre
```

Important Considerations

- The software determines whether there is a better path for the frame by determining whether the source interface is the same as the destination interface and whether the frame's sender is on a directly connected network. If the software determines that a received frame has a better path available through another gateway:
 - It sends an ICMP redirect message back to the originator of the frame indicating the better gateway to use in the future
 - It routes the frame to the gateway
- ICMP redirect can be set on a per-interface basis.
- For better performance or if you have applications that ignore ICMP redirects, disable the ability of the interface to send ICMP redirects.
- If you have two interfaces that belong to virtual LANs (VLANs) that share a given port and you want to completely disable ICMP redirects for that port, disable the redirects for each interface that shares that port. If you disable it for only one interface and enable it for the other, you may not get the performance improvement that you want.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces to which you want to enable or disable the transmission of an ICMP redirect to the sender of a frame. (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Current value
ICMP redirect state	Whether you want to implement ICMP redirect state on an interface. The system prompts you for a state for each interface.	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled, or current value

**ip interface
icmpRouterDiscovery**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables Internet Control Message Protocol (ICMP) router discovery, which enables hosts that are attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers and determine which router to use for a default gateway.

Valid Minimum Abbreviation

```
ip i icmpro
```

Important Considerations

- ICMP router discovery can be set on a per-interface basis.
- When you enable the state for an interface, the system prompts you for a preference. (See RFC 1256.) By default, this preference level is 0. Use the preference to control the use of certain routers as the default router. The host uses the router with the highest preference level.
- An appropriately configured end station can locate one or more routers on the LAN to which it is attached. The end station then automatically installs a default route to each of the routers that are running Internet Control Message Protocol (ICMP) router discovery. You do not need to manually configure a default route. ICMP redirect messages subsequently channel the IP traffic to the correct router.
- You can configure only certain end stations to work with the ICMP router discovery protocol. See the documentation for your workstation to determine whether you can configure it to work with this protocol.
- You can configure and display three timers for ICMP router discovery on the CoreBuilder 3500:
 - **Maximum advertisement interval** — The maximum time interval between advertisements.
 - **Minimum advertisement interval** — The minimum time interval between advertisements.
 - **Advertisement holdtime** — The length of time that advertisements are held valid.



The ranges for minimum advertisement interval depend on the set values for maximum advertisement interval and the holdtime range depends on the input values for both the maximum and minimum advertisement intervals.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to enable or disable ICMP router discovery. (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Current value
Router discovery state	Whether you want to implement ICMP router discovery on an interface. The system prompts you for a state for each interface.	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled, or current value
Preference	If you select <code>enabled</code> , the host interprets an unsigned integer as a 32-bit signed twos-complement integer that represents the preference level to associate with the interface. Higher values produce higher preference levels. The minimum value is reserved so that the address is not used as a default router address, only for specific IP destinations.	minimum value (hex 80000000) -2^{31} to 2^{31}	0
Maximum advertisement interval	Maximum interval between advertisements.	4 – 1800 seconds	600
Minimum advertisement interval	Minimum interval between advertisements.	3 – 600 seconds	450
Advertisement holdtime	Length of time that advertisements are held valid.	600 – 9000 seconds	1800

IP Interface ICMP Router Discovery Example (3500)

Select menu option (ip/interface): icmprouterdiscovery

Select IP interfaces (1|all|?) [1]: 1

Interface 1 - Enter router discovery state (disabled,enabled) [disabled]: enabled
Interface 1 - Enter router discovery preference [0]:
Interface 1 - Enter maximum advertisement interval (4-1800) [600]:
Interface 1 - Enter minimum advertisement interval (3-600) [450]:
Interface 1 - Enter advertisement holdtime (600-9000) [1800]:

ip interface statistics *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays IP interface statistics on a per-interface basis.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ip i st`

Important Consideration

- The system prompts you for an interface index number even if you have only one interface defined.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the interface whose statistics you want to display	<ul style="list-style-type: none"> ■ One or more configured interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	–

Fields in the IP Interface Statistics Display

Field	Description
inAddrErrors	Number of datagrams that the IP station discarded because of an error in the source or destination IP address
inCsumErrors	Number of datagrams that were dropped because of a checksum error
inDelivers	Number of datagrams that the IP station delivered to local IP client protocols
inDiscards	Number of packet receive discards
inForwards	Total number of packets that were forwarded (that is, routed through hardware or software or both)
inHdrErrors	Number of datagrams that the IP station discarded because the IP header contained errors
inReceives	Total number of IP datagrams that were received, including those with errors
inSameSegment	Number of packets that were received on an interface and that need to be forwarded out on the same interface

Field	Description
inTtlExceeds	Number of packets that were received on an interface and that need to be forwarded, but that have an IP header TTL value of less than 2
outDiscards	Number of packet transmit discards
outForwards	Total number of packets that a router has forwarded to an outbound interface (that is, routed through hardware or software or both)

ip route display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
✓ 9400

Displays the system's routing table to determine which routes to other IP networks are configured and whether the routes are operational.

Valid Minimum Abbreviation

```
ip route di
```

✓ 3900
✓ 9300

Important Considerations

- For the CoreBuilder 3500 only, the system prompts you for an IP address and subnet mask. As a result, you can display only a subset of routes instead of all routes. To see all entries in the table, simply press Enter at these prompts.
- The first line in the output (the status line) indicates whether IP routing is enabled:
 - For the CoreBuilder 9000 Layer 3 module, it also indicates whether Internet Control Message Protocol (ICMP) router discovery is enabled on the system.
 - For the CoreBuilder 3500, IP interface options (such as ICMP router discovery) appear under "ip interface detail" earlier in this chapter and are set on a per-interface basis.

Options (3500 only)

Prompt	Description	Possible Values	[Default]
IP address	IP address (and its corresponding subnet mask) for which to display only those routes that match the bits set in it	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 (displays all entries) 	0.0.0.0
Subnet mask	Subnet mask for the specified IP address for which to display only those routes that match the bits set in it	A valid subnet mask of a specified IP address	Current value

Fields in the IP Route Display

Field	Description
Destination	IP address of the destination network, subnetwork, or host. This field can also identify a default route, which the system uses to forward packets that do not match any other routing table entry. You may want to use the default route in place of routes to numerous destinations that all have the same gateway IP address.
Subnet mask	Subnet mask that is associated with the IP address of the destination network, subnetwork, or host.
Metric	Associated cost of sending a packet to the destination. The system includes the metric in its RIP and OSPF updates to allow other routers to compare routing information received from different sources.
Gateway	Address that directs the router how to forward packets whose destination addresses match the route's IP address and subnet mask. The system forwards such packets to the indicated gateway.
Status	Status of the route. See the following status table.
TTL	Time To Live — Time remaining before the route expires or is reset.

Status for Routes

Field	Description
Direct	Route is for a directly connected network
Learned	Route was learned using indicated protocol
Learned RIP-Zombie	Route was learned but is partially timed out. This condition is applied to all learned routes reached by an interface gateway which is in the down state.
Learned RIP2	Route was learned using RIP-2 protocol
Local	Actual interface address
Static	Route was statically configured
Timed out	Route has timed out and is no longer valid

ip route static *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a static route.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip route s`

✓ 3900

✓ 9300

Important Considerations

- Before you can define static routes, you must define at least one IP interface. See “ip interface define (3500/9000 Layer 3)” earlier in this chapter for more information.
- For the CoreBuilder 3500, you can define up to 256 static routes.
- For the other platforms, you can define up to 64 static routes.
- Static routes remain in the table; you must remove them before you can remove the corresponding interface.
- Static routes take precedence over dynamically learned routes to the same destination
- Static routes are included in periodic Routing Information Protocol (RIP) updates that the system sends.

Options

Prompt	Description	Possible Values	[Default]
Destination IP address	IP address of the destination network, subnetwork, or host for this route	A valid IP address	–
Subnet mask	Subnet mask for the specified IP address	A valid subnet mask	Based on specified IP address
Gateway IP address	IP address of the gateway that this route uses	A valid router address	–

ip route remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes an existing route.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip route r`

✓ 3900

✓ 9300

Important Consideration

- When you enter the command, the system deletes the route immediately from the routing table. You are not prompted to confirm the deletion.

Options

Prompt	Description	Possible Values	[Default]
Destination IP address	IP address of the route that you want to delete	A valid IP address	–
Subnet mask	Subnet mask for the specified IP address	A valid subnet mask	Based on specified IP address

ip route flush *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes all learned routes from the routing table.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

```
ip route fl
```

✓ 3900

✓ 9300

Important Considerations

- The system flushes all learned routes from the routing table immediately. You are not prompted to confirm the deletion.
- Flushing the routing table does not cause the Routing Information Protocol (RIP) to update the routing table. You must change the metric to update the routing table.

ip route default *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- ✓ 9400

Adds a default route to the routing table immediately.

Valid Minimum Abbreviation

`ip route de`

- ✓ 3900
- ✓ 9300

Important Considerations

- If you define a default route, the system uses it to forward packets that do not match any other routing table entry. The system can learn a route through the routing protocol, or you can statically configure a default route.
- The system can learn a default route.
- If the routing table does not contain a default route, the system cannot forward a packet that does not match any other routing table entry. When the system drops the packet, it sends an Internet Control Message Protocol (ICMP) `destination unreachable` message to the host that sent the packet.
- On the CoreBuilder 3500 or the CoreBuilder 9000 Layer 3 module, you establish a static sink default route, so that the system can advertise itself as a default router. The static sink default route is not used in any of the system's forwarding decisions because it does not have a valid next-hop gateway, but it can be advertised to all of the system's neighbors (unless you establish IP policies to prevent the advertisement).

Defining a static sink default route causes the route to be advertised through any IP protocols that you have configured on the system (for example, Open Shortest Path First (OSPF) and RIP). For more information about static sink default routes, see the *Implementation Guide* for the CoreBuilder 3500 or for the CoreBuilder 9000.

Options

Prompt	Description	Possible Values	[Default]
Gateway IP address	IP address of the route that you want to add as the default	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 (static sink default route) 	–

ip route noDefault *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes the default route.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip route n`

✓ 3900

✓ 9300

Important Consideration

- The system deletes the default route from the routing table immediately after you enter the command. You are not prompted to confirm the deletion.

ip route findRoute *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Searches for a route in the routing table.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

```
ip route fi
```

3900
9300

Important Considerations

- This command enables you to find a route using an IP address or a host name, as long as the Domain Name System (DNS) is configured.
- When you enter this command with a valid IP address or host name, the system displays the routing table entry.

Options

Prompt	Description	Possible Values	[Default]
IP address (or host name)	IP address of the route that you want to find, or a host name, if DNS is configured	<ul style="list-style-type: none"> ■ A valid IP address ■ A valid host name 	0.0.0.0, or current value

ip arp display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Displays the contents of the Address Resolution Protocol (ARP) cache for each interface on the system.

Valid Minimum Abbreviation

```
ip ar d
```

Important Considerations

- The system uses the ARP cache to find the MAC addresses that correspond to the IP addresses of hosts and other routers on the same subnetworks. Each device that participates in routing maintains an *ARP cache*, which is a table of known IP addresses and their corresponding MAC addresses.
- The first line in the output (the status line) indicates whether IP routing is enabled:
 - For the CoreBuilder 9000 Layer 3 module, it also indicates whether Internet Control Message Protocol (ICMP) router discovery is enabled on the system.
 - For the CoreBuilder 3500, IP interface options (such as ICMP router discovery) appear under “ip interface detail” earlier in this chapter and are set on a per-interface basis. The second status line indicates the number of entries in the ARP cache.

Fields in the IP ARP Display

Field	Description
Circuit	Circuit identifier
Hardware address	MAC address that is mapped to the IP address
I/F	Index number of the associated interface
IP address	IP address of the interface
Type	Type of entry — <i>static</i> or <i>dynamic</i>

ip arp static *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- ✓ 9400

Defines a static Address Resolution Protocol (ARP) cache entry on the system.

Valid Minimum Abbreviation

ip ar s

- ✓ 3900
- ✓ 9300

Important Considerations

- For the CoreBuilder 3500, you can define up to 128 static ARP entries.
- For the other platforms, you can define up to 64 entries.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the interface for which you want to define a static ARP entry	<ul style="list-style-type: none"> ■ A selectable interface index ■ ? (for a list of selectable interface indexes) 	–
IP address	IP address to use in the entry	A valid IP address	–
MAC address	Hardware address to use in the entry	A valid MAC address in the format XX-XX-XX-XX-XX-XX	–

IP ARP Static Example

```
Select interface index {1-2|?} 2
Enter IP address: 158.101.12.12
Enter MAC address: 00-00-00-00-00-01
```

ip arp remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Deletes an entry from the Address Resolution Protocol (ARP) cache (for example, if the MAC address has changed).

Valid Minimum Abbreviation

```
ip ar rem
```

Important Considerations

- When you enter the command, the system deletes the entry from the cache immediately. You are not prompted to confirm the deletion.
- If necessary, the system subsequently uses ARP to find the new MAC address that corresponds to that IP address.

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address for the entry that you want to delete	A valid IP address	–

ip arp flushAll *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes all entries from the Address Resolution Protocol (ARP) cache.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip ar flushA`

✓ 3900

✓ 9300

Important Considerations

- To flush dynamic entries only, see “ip arp flushDynamic” next in this chapter.
- When you enter the command, the system deletes all entries from the cache immediately. You are not prompted to confirm the deletion.

ip arp flushDynamic *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

✓ 9400

Deletes all dynamic (learned) entries from the Address Resolution Protocol (ARP) cache.

Valid Minimum Abbreviation

```
ip ar flushD
```

✓ 3900

✓ 9300

Important Considerations

- To flush *all* entries, static and dynamic, see the previous “ip arp flushAll” option.
- When you enter the command, the system deletes all dynamic entries from the cache immediately. You are not prompted to confirm the deletion.

ip arp age *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Sets the age time for dynamic Address Resolution Protocol (ARP) cache entries.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

ip ar a

- ✓ 3900
- ✓ 9300

Important Considerations

- The *age time* determines how long, in minutes, that the dynamic entries remain in the ARP cache before they are removed.
- By default, the system flushes the entry from the cache when it reaches the age time.
- A value of 0 indicates no age time, and the entry remains in the table until you remove it with the `ip arp remove` option or flush the ARP cache with the appropriate `flush` option.

Options

Prompt	Description	Possible Values	[Default]
Age time	Time that dynamic entries remain in the ARP cache	<ul style="list-style-type: none"> ■ 1 – 1440 minutes ■ 0 (to disable aging) 	15 (factory default), or current value

ip arp statistics *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Displays detailed information about the specified interfaces or all interfaces.

Valid Minimum Abbreviation

`ip ar status`

Important Considerations

- Your system tracks the number of times that a particular Address Resolution Protocol (ARP) event occurs.
- If a port that has multiple IP interfaces associated with it receives an ARP frame that is discarded because of an address mismatch, the `inReceives` and `inDiscards` statistics are incremented for the *first* interface of all the interfaces that are associated with the port.
- The system supports baselining for ARP statistics.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the IP interface from which to select ARP statistics	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	1

Fields in the IP ARP Statistics Display

Field	Description
<code>inDiscards</code>	<p>Received ARP frames that have been discarded due to one of the following reasons:</p> <ul style="list-style-type: none"> ■ Frame had a source address that did not match any directly connected IP interface that was associated with the port on which it was received ■ Frame contained an invalid header ■ Frame was not an ARP request or an ARP reply
<code>inReceived</code>	ARP frames (requests, replies, and discards) that were received on an IP interface

Field	Description
inReplies	ARP reply frames that were received on an IP interface
inRequests	ARP request frames that were received on an IP interface
outIfdown	Failure of the system to send one of the following three frames because the state of the IP interface was down: <ul style="list-style-type: none">■ ARP request■ ARP reply■ IP frame to be forwarded (pending ARP resolution)
outMemErrors	Failure of the system to allocate memory to transmit either an ARP request or an ARP reply
outReplies	ARP replies that were transmitted from an IP interface
outRequests	ARP requests that were transmitted from an IP interface

ip dns display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Displays the current domain name and the name servers that are associated with it.

Valid Minimum Abbreviation

```
ip d di
```

Important Considerations

- The *Domain Name System (DNS)* client provides DNS lookup functionality to the CoreBuilder IP ping and traceRoute features. You can specify a host name rather than an IP address when you perform various operations (for example, when you use ping or traceRoute to contact an IP station).
- With the DNS commands, you specify one or more name servers that are associated with a domain name. Each name server maintains a list of IP addresses and their associated host names. When you use ping or traceRoute with a host name, the DNS client attempts to locate the name on the name servers that you specify. When the DNS client locates the name, it resolves it to the associated IP address.
- See UNIX Network File System (NFS) documentation for information about how to create and maintain lists of domain names and IP addresses on the name servers.

Fields in the IP DNS Display

Field	Description
Domain name	Name of the domain name (up to 79 alphanumeric characters)
Name server	Name server that is associated with the domain

ip dns domainName *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Changes the name of a currently defined domain.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

ip d do

- ✓ 3900
- ✓ 9300

Important Considerations

- You can specify a domain name with up to 79 alphanumeric characters.
- Use single quotation marks (' ') around any string that has embedded spaces. Use double quotation marks (" ") to enter an empty string.

Options

Prompt	Description	Possible Values	[Default]
Domain name	Name of the domain. The name can be up to 79 characters long.	<ul style="list-style-type: none"> ■ A valid domain name ■ ? (to get information about specifying a domain name) 	– (or current name)

ip dns define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

✓ 9400

Defines a new name server IP address to associate with the current domain name.

Valid Minimum Abbreviation

`ip d de`

✓ 3900

✓ 9300

Important Considerations

- When the system accepts the new IP address, it displays a message like the following:

Server's IP address `xxxxx` is added to the DNS database

- The system assigns an index number to the new IP address. Use this index number to modify or remove this IP address.

Options

Prompt	Description	Possible Values	[Default]
Name server IP address	IP address of the name server that you want to define	A valid IP address	–

ip dns modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Modifies a currently defined name server IP address.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

ip d m

- ✓ 3900
- ✓ 9300

Important Considerations

- When you enter the command, the system displays the list of name server addresses and the index number that is associated with each.
- The system assigns an index number to the new IP address. Use this index number to modify this IP address.

Options

Prompt	Description	Possible Values	[Default]
Index	Index number of the name server IP address that you want to modify	<ul style="list-style-type: none"> ■ A selectable server index number ■ ? (for a list of selectable server indexes) 	–
Name server IP address	New IP address of the name server that you want to use	A valid IP address	–

ip dns remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes a previously defined name server IP address.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip d r`

✓ 3900

✓ 9300

Important Consideration

- When you enter the command, the system displays the list of name server addresses and the index number that is associated with each.

Options

Prompt	Description	Possible Values	[Default]
Index	Index number of the name server IP address that you want to remove	<ul style="list-style-type: none"> ■ A selectable server index number ■ ? (for a list of selectable server indexes) 	–

ip dns nslookup *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- ✓ 9400

Maps an IP address to a host name or a host name to an IP address on a name server.

Valid Minimum Abbreviation

ip d n

- ✓ 3900
- ✓ 9300

Important Considerations

- Specify a host name or IP address at the prompt.
- Enter a string of up to 255 characters.
- Use single quotation marks (' ') around any string with embedded spaces. Use double quotation marks (" ") to enter an empty string.

Options

Prompt	Description	Possible Values	[Default]
IP address or host name	IP address or host name that you want to map	<ul style="list-style-type: none"> ■ A host name of up to 255 characters ■ A valid IP address 	–

ip udpHelper display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Displays the BOOTP (bootstrap protocol) hop count and the threshold configuration. Also lists the ports with their IP forwarding addresses that are defined in your system.

3900
9300

Valid Minimum Abbreviation

`ip u di`

Important Considerations

- With UDP Helper, you can send User Datagram Protocol (UDP) packets between routed networks. UDP Helper provides support for UDP services such as BOOTP and DHCP (Dynamic Host Configuration Protocol), which rely on the BOOTP relay agent.
- When you configure the logical BOOTP port, you can boot hosts through the router. UDP Helper also provides a relay agent for DHCP broadcasts. UDP packets that rely on the BOOTP relay agent are modified and then forwarded through the router.
- BOOTP (including DHCP) uses UDP port 67.
- With UDP Helper, you can configure the amount of time that a UDP packet is forwarded between subnetworks. The system discards UDP packets based on the hop count and the seconds value only for BOOTP and DHCP.

Fields in the IP udpHelper Display

Field	Description
UDP port	UDP port number — usually the value 67
Forwarding address	Forwarding address that is used for UDP packets

ip udpHelper define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines port numbers or IP forwarding addresses for the UDP Helper.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip u de

3900
9300

Important Considerations

- You can have up to 63 combinations of port numbers and IP forwarding addresses per router.
- You can have multiple IP address entries for the same ports.

Options

Prompt	Description	Possible Values	[Default]
UDP port number	Port number for UDP	1 – 65535	67 (factory default), or current value
IP forwarding address	Forwarding addresses that are used for UDP packets	A valid IP address	–

ip udpHelper remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Removes a port number or IP forwarding address that has been defined for UDP Helper.

Valid Minimum Abbreviation

ip u r

Important Consideration

- The system immediately removes the port numbers and IP forwarding addresses that you specified. You are not prompted to confirm the deletion.

Options

Prompt	Description	Possible Values	[Default]
UDP port number	UDP port number that you want to remove	1 – 65535	67 (factory default), or current value
IP forwarding address	Forwarding addresses that you want to remove	A valid IP address	–

**ip udpHelper
hopCountLimit*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Sets the maximum hop count to specify how many steps the system uses to forward a packet through the router.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**

ip u h

Options

Prompt	Description	Possible Values	[Default]
BOOTP hop count limit	Maximum number of hops to allow for UDP packet forwarding	0 – 16	4 (factory default), or current value

**ip udpHelper
threshold****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Sets the maximum number of times that the system forwards a packet to the network.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip u t

Important Consideration

- By default, there is no threshold (0).

Options

Prompt	Description	Possible Values	[Default]
BOOTP relay threshold	Maximum number of times that the system forwards a packet to the network	0 – 65535	0 (factory default), or current value

**ip udpHelper
interface first****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Configures UDP Helper to support overlapped IP interfaces by using the first interface.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**

ip u i f

3900
9300**Important Considerations**

- *Overlapped* IP interfaces are multiple logical interfaces that are defined for a single physical port. You can specify how UDP Helper forwards packets from overlapped IP interfaces with one of three interface options (*first*, *even*, or *sequential*).
- The value `first` directs the system to use the first overlapped IP interface as the source network for forwarded packets.
- The system implements your selection immediately. You can view the UDP Helper configuration when you configure the forwarding address.

**ip udpHelper
interface even**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Configures UDP Helper to support overlapped IP interfaces by evenly distributing interfaces.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`ip u i e`

3900
9300

Important Considerations

- The value `even` directs the system to hash the client's MAC address to determine the source network for forwarded packets. This arrangement evenly distributes the interface among those on the network.
- The system implements your selection immediately. You can view the UDP Helper configuration when you configure the forwarding address.

**ip udpHelper
interface sequential**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Configures UDP Helper to support overlapped IP interfaces by distributing the interfaces sequentially.

Valid Minimum Abbreviation

`ip u i s`

Important Considerations

- The value `sequential` directs the system to assign each overlapped IP interface, in turn, as the source network for forwarded packets.
- The system implements your selection immediately. You can view the UDP Helper configuration when you configure the forwarding address.

ip routing *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Controls whether the system forwards or discards IP packets that are addressed to other hosts.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

```
ip routi
```

Important Considerations

- When you enable IP routing, the system acts as a standard IP router: it forwards IP packets from one subnetwork to another when required.
- When you disable IP routing, the system discards any IP packets that are not addressed directly to one of its defined IP interfaces.
- By default, IP routing is `disabled` on the system.

Options

Prompt	Description	Possible Values	[Default]
IP routing state	Whether IP routing is implemented on the system	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled

ip rip display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
✓ 9400

Displays information about the Routing Information Protocol (RIP) interfaces on the system. RIP is one of the IP Interior Gateway Protocols (IGPs). When RIP is enabled, the system dynamically configures its routing tables.

✓ 3900
✓ 9300

Valid Minimum Abbreviation

`ip ri d`

Important Considerations

- The output for this display differs according to platform.
- The first line in the output (the status line) indicates whether IP routing is enabled:
 - For the CoreBuilder 9000 Layer 3 module, it also indicates whether Internet Control Message Protocol (ICMP) router discovery is enabled on the system.
 - For the CoreBuilder 3500, IP interface options (such as ICMP router discovery) appear under “ip interface detail” earlier in this chapter and are set on a per-interface basis. The rest of the output contains more RIP interface information.
- The four available RIP modes are as follows:
 - **Disabled** — The system ignores all incoming RIP packets and does not generate any RIP packets of its own.
 - **Learn** — The system processes all incoming RIP packets, but it does not transmit RIP updates.
 - **Advertise** (3500 and 9000 only) — The system broadcasts RIP updates, but it does not process incoming RIP packets.
 - **Enabled** (3500 and 9000 only) — The systems broadcasts RIP updates and processes incoming RIP packets.
- An advertising router sends a RIP message every 30 seconds with both the IP address and a *metric* (the distance to the destination from that router) for each destination. Each router through which a RIP packet must travel to reach a destination equals one *hop*.

Fields in the IP RIP Display

Field	Description
Advertisement Addresses (3500 and 9000 only)	List of available advertisement addresses. The list is used for RIP-2 updates only if the RIP-1 compatibility mode is enabled. RIP-1 always uses advertisement addresses.
Compatibility Mode (3500 only)	Whether RIP 1 compatibility mode is <code>enabled</code> or <code>disabled</code> (by default, <code>disabled</code>).
Cost (3500 and 9000 only)	RIP cost for the interface (by default, 1).
Index	Index number of the interface.
Poison Reverse (3500 and 9000 only)	Whether poison reverse mode is <code>enabled</code> or <code>disabled</code> (by default, <code>enabled</code>).
RIP-1 Mode	Mode for RIP-1. If you disable RIP-1, the output lists the state as <code>off</code> . Other modes are <code>learn</code> (default), <code>advertise</code> , and <code>enabled</code> .
RIP-2 Mode	Mode for RIP-2. If you disable RIP-2, the output lists the state as <code>off</code> . Other modes are <code>learn</code> (default), <code>advertise</code> , and <code>enabled</code> .
Route Aggregate (3500 only)	Whether Route Aggregation mode is <code>enabled</code> or <code>disabled</code> .

ip rip mode *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
✓ 9400

On a per-interface basis, sets one of four RIP Version 1 (RIP-1) modes on the system. For all platforms except the CoreBuilder 9000, also allows you to set RIP Version 2 (RIP-2) modes.

Valid Minimum Abbreviation

ip ri m

✓ 3900
✓ 9300

Important Considerations

- Platforms except the CoreBuilder 9000 support RIP Version 1 as well as RIP Version 2. For each interface, you select a RIP Version 1 mode and a RIP Version 2 mode. The default RIP Version 1 mode for all platforms is `learn`. The default RIP Version 2 mode is `learn`.
- The four available RIP modes are as follows:
 - **Disabled** — The interface ignores all incoming RIP packets and does not generate any RIP packets of its own.
 - **Learn** — The interface processes all incoming RIP packets, but it does not transmit RIP updates. This is the default RIP mode.
 - **Advertise** (3500 and 9000 only) — The interface broadcasts RIP updates, but it does not process incoming RIP packets.
 - **Enabled** (3500 and 9000 only) — The interface broadcasts RIP updates and processes incoming RIP packets.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the RIP mode	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable

Prompt	Description	Possible Values	[Default]
RIP mode, Version 1	Selected RIP Version 1 mode that determines how the interface handles RIP 1 packets and updates	<ul style="list-style-type: none"> ■ disabled ■ learn ■ advertise (3500/9000) ■ enabled (3500/9000) 	learn (factory default), or current value
RIP mode, Version 2 (not 9000)	Selected RIP mode that determines how the interface handles RIP 2 packets and updates	<ul style="list-style-type: none"> ■ disabled ■ learn ■ advertise (3500 only) ■ enabled (3500 only) 	learn (factory default), or current value

IP RIP Mode Example

```
Select IP interfaces (1,2|all|?): 1
```

```
Interface 1 - Enter RIP Version 1 mode  
(disabled,learn,advertise,enabled) [learn]: disabled
```

```
Interface 1 - Enter RIP Version 2 mode  
(disabled,learn,advertise,enabled) [learn]: enabled
```

ip rip
compatibilityMode

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

On a per-interface basis, sets the RIP Version 1 compatibility mode.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip ri com

Important Considerations

3900
9300

- The RIP-1 compatibility mode determines how the software sends periodic RIP-2 updates. (For RIP-1, the software never uses the multicast address; it uses the advertisement list.)
 - When the system is configured to advertise RIP-2 packets and compatibility mode is `disabled`, the software uses the multicast address of 224.0.0.9 when sending periodic updates. This latest industry recommendation reduces the load on hosts that are not configured to listen to RIP-2 messages.
 - When the system is configured to advertise RIP-2 packets and compatibility mode is `enabled`, the software uses the advertisement list for RIP-2 updates.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the RIP compatibility mode	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
RIP-1 compatibility mode	Selected RIP Version 1 compatibility mode that determines how the system handles RIP-2 updates	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled

ip rip cost *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

On a per-interface basis, sets the RIP cost.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ip ri cos`

Important Considerations

- The default cost value is 1, which is appropriate for most networks.
- The system uses the cost number, between 1 and 15, to calculate route metrics. Unless your network has special requirements, assign a cost of 1 to all interfaces.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the RIP cost	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
RIP cost	Selected RIP cost for the interface	1 – 15	1 (factory default), or current value

ip rip poisonReverse *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Enables or disables RIP Poison Reverse mode on the system.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ip ri poi`

Important Considerations

- Your system always implements *Split Horizon*, a scheme that aims to avoid the problems that are associated with reverse-route updates (that is, the updates that are sent to a neighboring router that include the routes that are learned from that router). The scheme omits the routes that are learned from one neighbor in the updates that are sent to that neighbor (the reverse routes). Poison reverse works with Split Horizon as follows:
 - When you enable *Poison Reverse* for use with the Split Horizon scheme (the default), the system advertises reverse routes in updates, but sets the metrics to 16 (infinity). Setting the metric to infinity breaks the loop immediately when two routers have routes that point to each other.
 - When you disable Poison Reverse for the Split Horizon scheme, reverse routes are simply not advertised.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the poison reverse mode	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
Poison Reverse mode	Whether you want to implement poison reverse for the selected interface	<ul style="list-style-type: none"> ■ disabled ■ enabled 	Current value

ip rip
routeAggregation
Mode

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the route aggregation mode.

Valid Minimum Abbreviation

ip ri ro

Important Considerations

- *Route aggregation mode* determines which route table entries are sent during a RIP Version 2 update.
 - If route aggregation mode is `enabled`, RIP-2 can function like RIP-1 and “collapse” route table entries for all subnets of a directly connected network. For example, if route aggregation is `enabled`, and the system is advertising subnets 150.100.31.0 and 150.100.32.0, only the entry for network 150.100.0.0 is sent in the update. With RIP Version 2, you *must* enable route aggregation mode if you want the interface to collapse the route table entries and function like RIP-1.
 - If route aggregation mode is `disabled` (the default), a RIP-2 update sends all routing table entries.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the route aggregation mode	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
Route aggregation mode	Whether you want to implement route aggregation on the selected interface	<ul style="list-style-type: none"> ■ disabled ■ enabled 	Current value

✓ 3500

✓ 9000

9400

3900

9300

ip rip password *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

3900
9300

Sets the IP RIP-2 password so that you can choose the IP interfaces that can put RIP-2 updates into their routing tables.

Valid Minimum Abbreviation

ip ri pa

Important Considerations

- If the sending interface has an IP RIP-2 password, the receiving interface must have the same IP RIP-2 password. If the receiving interface has a different password or a null password, its routing table is not updated.
- If you are using RIP-1, do not use the password option.
- You cannot use the ASCII string `none` as the password. This string is reserved to indicate the default password, which is a null value.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the IP interfaces that you want to allow to receive route updates	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	–
Password	Combination of characters that you set as the RIP-2 password	<ul style="list-style-type: none"> ■ up to 16 alphanumeric characters ■ null password 	null password

IP RIP Password Example

```
Select menu option (ip/rip): password
Select IP interfaces (1,2|all|?): 1
Interface 1 - Enter password {?} [none]: wings
```


ip rip
addAdvertisement

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Adds an advertisement address to an IP RIP interface.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip ri a

Important Considerations

- The system uses the specified advertisement address to advertise routes to other stations on the same network. It uses this address for sending updates. (RIP-2 updates depend on the setting of RIP compatibility mode.)
- Advertisement addresses are handled differently based on RIP-1 and RIP-2.
 - For the CoreBuilder 3500, each interface that you define initially uses the default broadcast address (255.255.255.255) as the advertisement address. With RIP-1 updates, the address that you specify becomes the new RIP-1 advertisement address if you change the broadcast address. If you subsequently use RIP-2 (configure the interface to send RIP-2 advertisements) and have the RIP-1 compatibility mode disabled, the multicast address is used for updates.
 - For the CoreBuilder 9000, each interface that you define initially uses the directed broadcast address as the RIP advertisement address (all 1s in the host field).
- You can specify up to 64 advertisement addresses in separate iterations.
- On the CoreBuilder 3500:
 - After you add an advertisement address, you cannot subsequently change the broadcast address.
 - If you are using RIP-2 for the interface, you must enable RIP compatibility mode if you want the system to use the advertisement list instead of the multicast address for RIP updates. See “ip rip compatibilityMode” earlier in this chapter for more information.
- To add an advertisement address on other platforms, you must remove the directed broadcast address if you only want the address that you added to be used for RIP advertisements.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to add the advertisement address	<ul style="list-style-type: none">■ One or more selectable interface indexes■ ? (for a list of selectable interface indexes)	Previous entry, if applicable
Advertisement address	Selected IP address to add to the list of advertisement addresses	A valid IP address	—

**ip rip remove
Advertisement*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Removes an advertisement address from the list of RIP advertisement addresses for an interface.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ip ri re`**Options**

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to remove the advertisement address	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
Advertisement address	Advertisement address that you want to remove	An address from the advertisement list	—

ip rip policy summary *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Displays summary information about RIP routing policies.

Valid Minimum Abbreviation

```
ip ri pol s
```

Important Considerations

- Your system has one unified IP routing table. Route policies enable you to control the flow of routing information between the network, the protocols, and the unified routing table on your system.
- Route policies are classified as follows:
 - *Import policies* import routing information from what RIP learns from a router/neighbor to the unified routing table. (You can also import routing information from Open Shortest Path First (OSPF).)
 - *Export policies* send information from the routing table to RIP and RIP routers, which controls what is going out on the wire to the RIP domain. (You can also export from the routing table to OSPF.)
- The system tracks policies that you define in both RIP and OSPF, so the indexes that are assigned to your policies may have gaps (for example, if you have RIP policies 1 and 2 and OSPF policies 3-6, the next policy that is available for RIP or OSPF is 7).

Fields in the IP RIP Policy Summary Display

Field	Description
Action	Action for the route — accept or reject
Index	Index number of the policy
Protocol	Protocol (for example, RIP)
Route	Route affects the policy
Source	Source router (a11 is from all routers)
Type	Whether the policy is an import or export policy
Weight	Administrative weight — 1 through 16

- ip rip policy detail** *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*
Displays detailed information about RIP routing policies.
- ✓ 3500
✓ 9000
9400
- Valid Minimum Abbreviation**
ip ri pol det
- 3900
9300
- Important Considerations**
- This display contains the summary information and two additional fields: Interface and Metric.
 - Route policies are classified as follows:
 - *Import policies* import routing information from what RIP learns from a router/neighbor to the unified routing table. (You can also import routing information from Open Shortest Path First (OSPF).)
 - *Export policies* send information from the routing table to RIP and RIP routers, which controls what is going out on the wire to the RIP domain. (You can also export from the routing table to OSPF.)

Fields in the IP RIP Policy Detail Display

Field	Description
Action	Action for the route — accept or reject
Index	Index number of the policy
Interface	Interface that is associated with the policy (all applies to all interfaces)
Metric	Assigned metric, a value 0 through 16 for RIP-1 or RIP-2 (metrics can use options +, -, /, *, and %)
Protocol	Protocol (for example, RIP)
Route	Route that the policy affects
Source	Source router (all is from all routers)
Type	Whether the policy is an import or export policy
Weight	Administrative weight — 1 through 16

ip rip policy define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines an import or export route policy for RIP.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

```
ip ri pol def
```

3900
9300

Important Considerations

- Route policies are classified as follows:
 - *Import policies* import routing information from what RIP learns from a router/neighbor to the unified routing table. (You can also import routing information from Open Shortest Path First (OSPF).)
 - *Export policies* send information from the routing table to RIP and RIP routers, which controls what is going out on the wire to the RIP domain. (You can also export from the routing table to OSPF.)
- The system assigns an index number to each policy and takes into account all route policies set on the system, RIP and OSPF (You can define up to 128 routing policies total, shared between OSPF and RIP policies).
- Certain conditions are associated with import and export policies. See the import and export policy tables that follow the Options table for lists of the conditions.
- You can set up an IP RIP or OSPF import or export policy to accept or advertise the default route, as long as the default route exists in the routing table. When you define a policy, you are always prompted for the route subnet mask after the route address, regardless of whether you specify the wildcard route address of 0.0.0.0. For more information about the default route and routing policies, see the *CoreBuilder 3500 Implementation Guide* or the *CoreBuilder 9000 Implementation Guide*.

Options

Prompt	Description	Possible Values	[Default]
Policy type	Type of policy	<ul style="list-style-type: none"> ■ import ■ export 	import
Origin protocols	Which protocol advertises the route (for export policies only)	<ul style="list-style-type: none"> ■ directory ■ static ■ rip ■ ospf ■ all 	static
Source address	Router's IP address	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 ■ all 	0.0.0.0
Route address	Associated route IP address	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 ■ all 	0.0.0.0
Route subnet mask	Subnet mask for the route (for example, 255.255.0.0)	A valid mask	Based on route
IP interfaces	Index number of the interface indexes for which you want to define a routing policy	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	all, or previous entry, if applicable
Policy action	Whether to accept or reject the route	<ul style="list-style-type: none"> ■ accept ■ reject 	accept
Metric adjustment	For accept conditions only, increase or decrease in the converted route metric by the specified value. Options: + (add) - (subtract) * (multiple metric by value) / (use new metric as divisor) % (modulus, remainder of division operation as integer)	0 – 16, with or without options	0, which does not change the metric

Prompt	Description	Possible Values	[Default]
Administrative weight	Metric value for this policy (higher values have higher priority)	1 – 16	1

RIP Import Policy Conditions for Specified Interfaces

Source Router	Route (address/mask)	Action	Description
Specified router	Specified route/mask	accept	Accept specified route from specified source router on specified interfaces with or without metric adjustments (+, -, *, /, %).
Specified router	all (0.0.0.0)	accept	Accept all routes from specified router on specified interfaces with or without metric adjustments (+, -, *, /, %).
all (all routers)	Specified route/mask	accept	Accept specified route on specified interfaces with or without metric adjustments (+, -, *, /, %).
all	all	accept	Accept all routes on specified interfaces with or without metric adjustments (+, -, *, /, %).
Specified router	Specified route/mask	reject	Reject specified route from specified router on specified interfaces. (Metrics are not applicable.)
Specified router	all	reject	Reject all routes from specified router on specified interfaces.
all	Specified route/mask	reject	Reject specified route from all routers on specified interfaces.
all	all	reject	Reject all routes on specified interfaces.

RIP Export Policy Conditions for Specified Interfaces

Protocol	Source Router	Route	Action	Description
RIP, OSPF, static	Specified router or all routers	Specified route/mask	accept	Advertise RIP/OSPF/static specified route from specified source router on specified interfaces with or without metric adjustments (+, -, *, /, %).
RIP, OSPF, static	Specified router or all routers	all (0.0.0.0)	accept	Advertise all RIP/OSPF/static routes from specified router on specified interfaces with or without metric adjustments (+, -, *, /, %).
RIP, OSPF, static	Specified router or all routers	Specified route/mask	reject	Do not advertise the RIP/OSPF/static specified route on specified interfaces.
RIP, OSPF, static	Specified routers or all routers	all	reject	Do not advertise all RIP/OSPF/static routes on specified interfaces.

Example of Import Policy

```
Select menu option (ip/rip/policy): define
Enter policy type (import,export) [import]: import
Enter source address [0.0.0.0]:
Enter route address [0.0.0.0]: 158.101.135.40
Enter route subnet mask [255.255.0.0]:
Select IP interfaces (1,2|all|?) [1]: 1
Enter policy action (accept,reject) [accept]:
Enter metric adjustment ([+, -, *, /, %] 0-16) [0]:
Enter administrative weight (1-16) [1]:
```

Example of Export Policy

```
Select menu option (ip/rip/policy): define
Enter policy type (import,export) [import]: export
Enter origin protocols (dir,sta,rip,ospf|all|?) : rip
Enter source address [0.0.0.0]:
Enter route address [0.0.0.0]:
Select IP interfaces (1,2|all|?) [1]: all
Enter policy action (accept,reject) [accept]:
Enter metric adjustment ([+, -, *, /, %] 0-16) [0]:
Enter administrative weight (1-16) [1]:
```

ip rip policy modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Modifies an existing route policy for RIP.

Valid Minimum Abbreviation

```
ip ri pol m
```

3900
9300

Important Considerations

- Route policies are classified as follows:
 - *Import policies* import routing information from what RIP learns from a router/neighbor to the unified routing table. (You can also import routing information from Open Shortest Path First (OSPF).)
 - *Export policies* send information from the routing table to RIP and RIP routers, which controls what is going out on the wire to the RIP domain. (You can also export from the routing table to OSPF.)
- The system assigns an index number to each policy that you define. This index takes into account all route policies set on the system, RIP and OSPF, so the assigned index can be higher than you may expect.

Options

Prompt	Description	Possible Values	[Default]
Policy type	Type of policy	<ul style="list-style-type: none"> ■ import ■ export 	import
Origin protocols (export)	Whether or not the route is a static route (for export policies only)	<ul style="list-style-type: none"> ■ RIP ■ OSPF ■ all 	–
Source address	IP address of the source router	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 ■ all 	0.0.0.0
Route address	Route that is associated with the source network	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 ■ all 	0.0.0.0
Route subnet mask	Subnet mask that is associated with the route	A valid mask	Based on source network (for example, 255.255.0.0)

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interface for which you want to define a routing policy.	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
Policy action	Whether the route is accepted or rejected	<ul style="list-style-type: none"> ■ accept ■ reject 	accept
Metric adjustment	Used with accept, increase or decrease in the converted route metric by the specified value Options include: + (add) - (subtract) * (multiple metric by value) / (use new metric as divisor) % (modulus, take remainder of division operation expressed as an integer)	0 – 16	0, which does not change the metric
Administrative weight	Metric value for this policy (higher values have higher priority over lower-numbered values associated with the route)	1 – 16	1

ip rip policy remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes a previously defined route policy.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

```
ip rip pol r
```

3900

9300

Important Considerations

- The system assigns an index number to each policy that you define. This index takes into account all route policies that are set on the system, RIP and OSPF, so the assigned index can be higher than you may expect.
- When you remove a policy, the associated index is available for future use.

Options

Prompt	Description	Possible Values	[Default]
Policy index	Index number that is associated with the policy that you want to delete	<ul style="list-style-type: none"> ■ One or more selectable policy indexes ■ all ■ ? (for a list of selectable policy indexes) 	–

ip rip statistics *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays general RIP statistics.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip rip s`

✓ 3900

✓ 9300

Fields in the IP RIP Statistics Display

Field	Description
queries	Number of queries
routeChanges	Number of route changes

ip ping *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Tries to reach or “ping” a specified destination using the default ping options.

✓ 3500
✓ 9000
✓ 9400

Valid Minimum Abbreviation

ip p

✓ 3900
✓ 9300

Important Considerations

- This tool is useful for network testing, performance measurement, and management. It uses the ICMP echo facility to send Internet Control Message Protocol (ICMP) echo request packets to the IP destination that you specify.
- If you need to change the default ping options, use the `ip advancedPing` option. (The command description for `ip advancedPing` lists the default ping options.)
- You can either supply the host name or IP address as part of the command string, or you can supply the information at the prompt.
- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “`ip dns domainName`” earlier in this chapter for more information.
- When the system sends an echo request packet to an IP station using ping, the system waits for an ICMP echo reply packet. Possible responses:
 - If the host is reachable, the system displays information about the ICMP reply packets and the response time to the ping.
 - If the host does not respond, the system displays the ICMP packet information and this message: `Host is Not Responding`. You may not have configured your gateway IP address.
 - If the packets cannot reach the host, the system displays the ICMP packet information and this message: `Host is Unreachable`. A host is unreachable when there is no route to that host.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping	<ul style="list-style-type: none"> ■ A valid host name ■ IP address 	0.0.0.0, or current value

IP Ping Example

```

Select menu option (ip): ping
Enter host name/IP address [0.0.0.0]: 158.101.111.50
Press "Enter" key to interrupt.

PING 158.101.111.50: 64 byte packets
64 bytes from 158.101.111.50: icmp_seq=0. time=16. ms
64 bytes from 158.101.111.50: icmp_seq=1. time=19. ms
64 bytes from 158.101.111.50: icmp_seq=2. time=24. ms

---- 158.101.111.50 PING Statistics ----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 16/20/24

```

ip advancedPing *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Tries to contact a host with one or more of the advanced ping options.

Valid Minimum Abbreviation

ip advancedP

Important Considerations

- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “ip dns domainName” earlier in this chapter for more information.
- The `burst` option, when enabled, overrides the value set in the `quiet` or `wait` option.
- The `burst` option floods the network with Internet Control Message Protocol (ICMP) echo packets and can cause network congestion. Do *not* use the `burst` option during periods of heavy network traffic. Use this option only as a diagnostic tool in a network that has many routers to determine if one of the routers is not forwarding packets. For example, you can set a high count value (1000 packets), and then observe the run lights on the units: the run lights blink rapidly on routers that are forwarding packets successfully, but remain unlighted, or blink slowly, on routers that are not forwarding packets successfully.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping.	<ul style="list-style-type: none"> ■ A valid host name ■ IP address 	0.0.0.0
Number of ICMP Request packets	Number of ICMP echo request packets that are sent to ping a host. If the destination host does not respond after it is pinged by the number of packets that you specify, the system displays a <code>Host is Unreachable</code> or <code>Host is not Responding</code> message.	1 – 9999 packets	3

Prompt	Description	Possible Values	[Default]
Packet size	Number of bytes in each ICMP echo request packet. The packet size includes both the IP and the ICMP headers.	28 – 4096 bytes	64
Burst Transmit Ping mode	How rapidly to send out ICMP echo request packets. When <code>enabled</code> , sends out the ICMP echo request packets as rapidly as possible. The system displays a period (.) upon receiving an ICMP echo replay packet. Use this display to determine how many packets are being dropped during the burst. This is unique to the burst option.	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled
Quiet mode	How much packet information to display after a ping. When <code>enabled</code> , the system displays information about the number of packets that the system sent and received, any loss of packets, and the average time that it took a packet to travel to and from the host. When <code>disabled</code> , the system displays more detailed status information about each ICMP echo request packet.	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled
Time between sending each packet (wait)	Number of seconds that the system waits before it sends out successive ICMP echo request packets. Set this option to a high value if network traffic is heavy and you choose not to add to the network traffic with pings in fast succession.	1 – 20 seconds	1
ICMP sourceAddress	Whether to force the source address of the ICMP packets to be something other than the IP address of the interface from which the packet originated. You can use this option if you have more than one IP interface defined.	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Interface index	Index number of the ICMP source IP address that you want to use. The system lists currently defined interfaces and their indexes.	A selectable interface index	0 (the router picks the best interface)

IP Advanced Ping Example

```
Select menu option (ip): advancedPing
Enter host IP address [0.0.0.0]: 158.101.112.56
Enter number of ICMP request packets (1-9999) [3]:
Enter packet size (bytes) (28-4096) [64]:
Enter Burst Transmit Ping mode (disabled,enabled) [disabled]:
Enter Quiet mode (disabled,enabled) [disabled]:
Enter time (sec) waits between sending each packet (1-20) [1]: 2
Configure ICMP sourceAddress? (n,y) [y]:
      Index      Interface address
      0          Best interface (default)
      1          158.101.117.151
      2          158.101.10.1
Select interface index {0-2|?} [0]: 1
Press "Enter" key to interrupt.

PING 158.101.112.56 from 158.101.117.151: 64 byte packets
64 bytes from 158.101.112.56: icmp_seq=0. time=26. ms
64 bytes from 158.101.112.56: icmp_seq=1. time=18. ms
64 bytes from 158.101.112.56: icmp_seq=2. time=18. ms

---- 158.101.112.56 PING Statistics ----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 18/21/26
```

ip traceRoute *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Traces a route to a destination using the default traceRoute options.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

ip t

✓ 3900

✓ 9300

Important Considerations

- TraceRoute information includes all of the nodes in the network through which a packet passes to get from its origin to its destination. It uses the IP time-to-live (TTL) field in UDP probe packets to elicit an Internet Control Message Protocol (ICMP) Time Exceeded message from each gateway to a host.
- To change the default traceRoute options, use `ip advancedTraceRoute`. (The command description for “ip advancedTraceRoute” lists the default traceRoute options.)
- You can either supply the host name or IP address as part of the command string, or you can supply the information at the prompt.
- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “ip dns domainName” earlier in this chapter for more information.
- To track the route of an IP packet, traceRoute launches User Datagram Protocol (UDP) probe packets with a small TTL value and then listens for an ICMP Time Exceeded reply from a gateway. Probes start with a small TTL of 1 and increase the value by 1 until one of the following events occurs:
 - The system receives a `Port Unreachable` message, which indicates that the packet reached the host.
 - The probe exceeds the maximum number of hops (default 30).

- At each TTL setting, the system launches three UDP probe packets, and the traceRoute display shows a line with the TTL value, the address of the gateway, and the round-trip time of each probe. If a probe answers from different gateways, the traceRoute feature prints the address of each responding system. If no response occurs in the 3-second timeout interval, traceRoute displays an asterisk (*) for that probe.

Other characters that can be displayed include the following:

- !N — Network is unreachable
 - !H — Host is unreachable
 - !P — Protocol is unreachable
 - !F — Fragmentation is needed
 - !<n> — Unknown packet type
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination to which you want to trace a route	<ul style="list-style-type: none"> ■ A valid host name ■ IP address 	0.0.0.0

IP Trace Route Example

```
Select menu option (ip): traceRoute
Enter host name/IP address [0.0.0.0]: 158.101.101.40
Press "Enter" key to interrupt.
```

```
Traceroute to 158.101.101.40: 30 hops max, 28 bytes packet
```

```
 1 158.101.117.254  9 ms 22 ms 5 ms
 2 158.101.112.254  8 ms 22 ms 8 ms
 3 158.101.96.22   7 ms 22 ms 7 ms
 4 158.101.101.40  7 ms 23 ms 6 ms
```

ip advancedTraceRoute

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Traces a route to a host with one or more of the advanced traceRoute options.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

ip advancedT

- ✓ 3900
- ✓ 9300

Important Considerations

- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “ip dns domainName” earlier in this chapter for more information.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping.	<ul style="list-style-type: none"> ■ A valid host name ■ IP address 	0.0.0.0
Maximum ttl	Maximum number of hops that the system can use in outgoing probe packets.	1 – 255 hops	30
Destination port	Destination (or base) UDP port number that the system uses in probe packets. Set the destination UDP port number to be very high to ensure that an application at the destination is not using that port.	30000 – 65535	33434
Probe count	Maximum number of probes that the system sends at each TTL level.	1 – 10	3
Wait	Maximum amount of time that the system waits for a response to a probe.	1 – 10 seconds	3
Packet size	Number of bytes that the system sends in each UDP probe packet.	28 – 4096 bytes	28
Source address	Source address other than the one from which the probe packets originate. This option is available if you have more than one IP interface defined on the system.	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the ICMP source IP address that you want to use The system lists defined interfaces and their indexes	A selectable interface index	0 (the router picks the best interface)
Numeric mode	Whether the system shows hop addresses numerically or symbolically	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled

IP Advanced Trace Route Example (TTL value of 10):

```

Select menu option (ip): advancedTraceRoute
Enter host IP address [158.101.101.27]:
Enter maximum Time-to-Live (ttl) (1-255) [30]: 10
Enter Destination Port number (30000-65535) [33434]:
Enter the number of probes to be sent at each ttl level (1-10) [3]:
Enter time (sec) to wait for a response (1-10) [3]:
Enter the packet size (bytes) (28-4096) [28]:
Configure TRACEROUTE sourceAddress? (n,y) [y]:
      Index      Interface address
      0          Best interface (default)
      1          158.101.117.151
      2          158.101.10.1
Select interface index {0-2|?} [0]:
Enter Numeric mode (disabled,enabled) [disabled]:
Press "Enter" key to interrupt.

```

Traceroute to 158.101.101.27: 10 hops max, 28 bytes packet

```

1  158.101.117.254  12 ms   7 ms   5 ms
2  158.101.112.254  51 ms   9 ms   7 ms
3  158.101.96.22   21 ms  15 ms   6 ms
4  158.101.101.27  18 ms  90 ms  80 ms

```

ip statistics *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Displays different types of IP statistics: general statistics and those specific to the User Datagram Protocol (UDP) or the Internet Control Message Protocol (ICMP).

Valid Minimum Abbreviation

ip sta

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Statistics	Type of IP statistics that you want to display	<ul style="list-style-type: none"> ■ ip ■ udp ■ icmp ■ all 	ip

Fields in the IP Statistics Display

Field	Description
forwDatagrams	Number of datagrams that the IP station tried to forward
fragCreates	Number of IP datagram fragments that were generated as a result of fragmentation on this system
fragFails	Number of ip datagrams that were discarded because they needed to be fragmented but could not be (for example, because their Don't Fragment bit was set)
fragOks	Number of IP datagrams that were successfully fragmented
inAddrErrors	Number of datagrams that the IP station discarded because of an error in the source or destination IP address
inDelivers	Number of datagrams that the IP station delivered to local IP client protocols
inDiscards	Number of packet receive discards
inHdrErrors	Number of datagrams that the IP station discarded because the IP header contained errors
inReceived	Total number of IP datagrams that were received, including those with errors
osReceives	Number of packets that were received that are destined to higher-level protocols such as Telnet, DNS, TFTP, and FTP
osTransmits	Number of packets that were sent through the router by higher-level protocols such as Telnet, DNS, TFTP, and FTP
outDiscards	Number of packet transmit discards

Field	Description
outNoRoutes	Number of datagrams that the IP station discarded because there was no route to the destination
outRequests	Number of datagrams that local IP client protocols passed to IP for transmission
reasmFails	Number of packet reassembly failures
reasmReqs	Number of packet reassembly requests
reasmOks	Number of successful packet reassemblies
rtDiscards	Number of packets that were discarded due to system resource errors
unkProtos	Number of packets whose protocol is unknown

Fields in the UDP Statistics Display

Field	Description
inDatagrams	Number of UDP packets that were received and addressed to the router or broadcast address
inErrors	Number of received UDP packets that contain header errors
noPorts	Number of UDP packets that were received but addressed to an unsupported UDP port
outDatagrams	Number of UDP packets that the router sent

Fields in the ICMP Statistics Display

Field	Description
inAddrMaskReps	Number of ICMP address mask reply frames that were received
inAddrMasks	Number of ICMP address mask request packets that were received
inDestUnreach	Number of ICMP destination unreachable packets that were received
inErrors	Number of received ICMP packets that contain header errors
inEchoReps	Number of ICMP echo reply packets that were received
inEchos	Number of ICMP echo request packets that were received
inParmProbs	Number of ICMP parameter problem frames that were received
inRedirects	Number of ICMP redirect packets that were received
inSrcQuenchs	Number of ICMP source quench packets that were received

Field	Description
inTimeExcds	Number of ICMP time exceeded packets that were received
inTimeStamps	Number of ICMP time stamp request packets that were received
inTimeStampsReps	Number of ICMP time stamp reply packets
messages	Number of ICMP packets that were received
outAddrMaskReps	Number of ICMP address mask reply packets that were sent
outAddrMasks	Number of ICMP address mask request packets that were sent
outDestUnreach	Number of ICMP destination unreachable packets that were sent
outEchoReps	Number of ICMP echo reply packets that were sent
outEchos	Number of ICMP echo request packets that were sent
outErrors	Number of ICMP packets that were sent that were dropped due to system resource errors
outMsgs	Number of ICMP packets that were sent
outParmProbs	Number of ICMP parameter problem packets that were sent
outRedirects	Number of ICMP redirect packets that were sent
outSrcQuenchs	Number of ICMP source quench packets that were sent
outTimeExcds	Number of ICMP time exceeded packets that were sent
outTimeStampReps	Number of ICMP time stamp reply packets that were sent
outTimeStamps	Number of ICMP time stamp request packets that were sent

17

VIRTUAL ROUTER REDUNDANCY (VRRP)

Virtual Router Redundancy Protocol (VRRP) provides fault-tolerant routing on a LAN by eliminating the single point of failure that exists when hosts are configured with a static default gateway. This chapter provides guidelines and other key information about configuring VRRP on your system.




For more information about VRRP, see the Implementation Guide for your system.



For the CoreBuilder® 9000 platform, the commands in this chapter apply to Layer 3 switching modules only.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



ip vrrp summary *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Displays summary information about configured virtual routers on your system.

Valid Minimum Abbreviation

```
ip v s
```

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) for which you want to display virtual router information	<ul style="list-style-type: none"> ■ One or more valid IP VLAN index numbers ■ all ■ ? (for a list of selectable indexes) 	—
Virtual router ID	ID of the virtual router for which you want to display summary information	<ul style="list-style-type: none"> ■ Valid virtual router ID (1 – 255) ■ ? (for list of selectable IDs) 	ID of virtual router that is defined on the VLAN

Fields in the IP VRRP Summary Display

Field	Description
Address	IP address of the virtual router
Auth	Whether the VRRP router uses simple password authentication. If password authentication is configured, the VRRP router discards any VRRP packet that does <i>not</i> have a matching authentication string.
Error	Last type of invalid advertisement received, or <i>none</i> .
Interval	Time, in seconds, between virtual router advertisements. The Master router advertises all IP addresses that are associated with the virtual router. Backup routers on the VRID consider the Master down if two advertisement intervals pass with no advertisement from the Master.
Ports	Ports that are defined on the virtual LAN (VLAN) and that are associated with the virtual router
Preempt	Whether a backup virtual router preempts a Master with a lower priority. <i>Yes</i> allows preemption; <i>no</i> prohibits it.

Field	Description
Pri	Priority of the the virtual router. Represented by a value from 0 through 255. Used in Master router election. Value of 255 indicates that the router owns the IP addresses that are associated with the virtual router. 0 indicates that the current Master has stopped participating in VRRP.
State	<p>Current state of the VRRP router. One of the following:</p> <ul style="list-style-type: none"> ■ Master — In this state, the router is the active forwarding router for all IP addresses that are associated with the virtual router. ■ Backup — In this state, the router monitors the availability of the Master router. If the Master router fails, the Backup router assumes forwarding responsibility for all IP addresses that are associated with the virtual router. ■ Initialize — Transitional state between Backup and Master states. Typically indicates that the virtual router has been configured but not enabled, or that the virtual router mode has been set to disabled. In this state, the router waits for a Startup event. When the router receives the Startup event, it broadcasts an ARP request that contains the virtual router MAC address for all IP addresses that are associated with the virtual router and transitions to the Master state. If the Startup event is not received, it transitions to the Backup state.
Type	Type of virtual router: primary or backup
VLAN Index	Index number of the virtual LAN (VLAN) on which the virtual router is defined
VRID	Virtual Router ID (0 – 255) . Must be unique on the LAN

Sample IP VRRP Summary Display

```
Select menu option (ip/vrrp): summary
Enter VLAN interface index (2|?) [2]:
Enter virtual router ID (1|?) [1]:
```

```
VLAN Index: 2 Ports: 7-12,14
```

VRID	Address	Type	State	Interval	Pri	Preempt	Auth	Error
1	158.101.175.228	Primary	Master	1 sec.	255	Yes	pass	none

ip vrrp detail *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Displays summary information and detailed statistics for the specified virtual router.

Valid Minimum Abbreviation

```
ip v det
```

Important Consideration

- Displays both summary information and the VRRP router statistics table for locally configured virtual routers, whether they are in the Master, Backup, or Initialize state.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) for which you want to display virtual router information	<ul style="list-style-type: none"> ■ One or more valid IP VLAN index numbers ■ all ■ ? (for a list of selectable indexes) 	—
Virtual router ID	ID of the virtual router for which you want to display summary information	<ul style="list-style-type: none"> ■ Valid virtual router ID (0 – 255) ■ ? (for list of selectable IDs) 	ID of virtual router that is defined on the VLAN

Fields in the IP VRRP Detail Display

Field	Description
Address	IP address of the virtual router
addrListErrors	Total number of VRRP advertisements that were received that do not match the address list defined for the virtual router
advertReceived	Total number of VRRP advertisements that this virtual router has received
advIntErrors	Total number of VRRP advertisement packets that were received for which the advertisement interval is different than the one that is configured for the virtual router

Field	Description
Auth	Whether the VRRP router uses simple password authentication. If password authentication is configured, the VRRP router discards any VRRP packet that does <i>not</i> have a matching authentication string.
authFailures	Total number of VRRP advertisements that this virtual router has received that did not have the correct simple text authentication password
becomeMaster	Total number of times that this virtual router has changed to the Master state
Error	Last type of invalid advertisement received, or <i>none</i> .
Interval	Time, in seconds, between virtual router advertisements. The Master router advertises all IP addresses that are associated with the virtual router. Backup routers on the VLAN consider the Master down if two advertisement intervals pass with no advertisement from the Master.
InvalidAuthType	Total number of VRRP advertisements that the virtual router has received with the Authentication Type not equal to the locally configured authentication method
invalidPktTypeRx	Number of VRRP advertisements with an invalid value in the Type field that this virtual router has received
ipTtlErrors	Total number of VRRP advertisements with IP TTL (Time-to-Live) not equal to 255 that this virtual router has received
MasterIpAdd	IP address of the Master for this virtual router.
Ports	Ports that are defined on the virtual LAN (VLAN) and that are associated with the virtual router
Preempt	Whether the router preempts a Master with a lower priority. <i>Yes</i> allows preemption; <i>no</i> prohibits it.
Pri	Priority of the virtual router. Represented by a value from 0 through 255. Used in Master router election. Value of 255 indicates that the router owns the IP addresses that are associated with the virtual router. 0 indicates that the current Master has stopped participating in VRRP.
PrimaryIpAddr	IP address which VRRP advertisements use as the source of the IP packet.
priorityZeroRx	Total number of VRRP advertisements with a priority of 0 that this virtual router has received. The priority of zero (0) indicates that the current Master has stopped participating in VRRP. Used to trigger Backup routers to quickly transition to Master without having to wait for the current Master to time out.

Field	Description
priorityZeroTx	Total number of VRRP advertisements with a priority of 0 that this virtual router has sent. The priority of zero (0) indicates that this virtual router was acting as Master but stopped participating in VRRP. Used to trigger backup routers to quickly transition to Master without having to wait for the current Master to time out.
State	<p>Current state of the VRRP router. One of the following:</p> <ul style="list-style-type: none"> ■ Master — In this state, the router is the active forwarding router for all IP addresses that are associated with the virtual router. ■ Backup — In this state, the router monitors the availability of the Master router. If the Master router fails, the Backup router assumes forwarding responsibility for all IP addresses that are associated with the virtual router. ■ Initialize — Transitional state between Backup and Master states. Typically indicates that the virtual router has been configured but not enabled, or that the virtual router mode has been set to disabled. <p>In this state, the router waits for a Startup event. When the router receives the Startup event, it broadcasts an ARP request that contains the virtual router MAC address for all IP addresses that are associated with the virtual router and transitions to the Master state. If the Startup event is not received, it transitions to the Backup state.</p>
Type	Type of virtual router: <code>primary</code> or <code>backup</code>
versionErrors	Total number of VRRP advertisements with an unknown or unsupported version number that this virtual router has received
VLAN Index	Index number of the virtual LAN (VLAN) on which the virtual router is defined
VRID	Virtual Router ID. Number that identifies the virtual router on the LAN

Sample IP VRRP Detail Display

```
Select menu option (ip/vrrp): detail
Enter VLAN interface index (2|?) [2]:
Enter virtual router ID (1|?) [1]:
```

```
VLAN Index: 2 Ports: 7-12,14
VRID Address          Type      State      Interval  Pri  Preempt  Auth  Error
  1  158.101.175.228  Primary Master  1 sec.  255   Yes   pass  none

VIDX VRID      becomeMaster  advertReceived  ckSumErrors  versionErrors
  2   1          1              0              0              0

VIDX VRID      advIntErrors  securViolations  ipTtlErrors  priorityZeroRx
  2   1          0              0              0              0

VIDX VRID      priorityZeroTx  invalidPktTypeRx  addrListErrors  unknownAuthType
  2   1          0              0              0              0

VIDX VRID      authTypeErrors
  2   1          0
```


ip vrrp define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a virtual router on the system.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

`ip v def`

3900

9300

Important Considerations

- Authentication passwords can be up to eight alphanumeric characters.
 - You can define *one* Primary router per VLAN.
 - Primary routers own the IP addresses that you associate with a virtual router.
 - When you define a Primary virtual router, the possible VLANs that you can select are the IP VLANs on the router that have no virtual routers configured.
 - The virtual router ID (VRID) must be unique across all locally attached LAN segments and unique for the local router.
 - When you define a Primary virtual router, you cannot use the VRID of a virtual router that is already defined on the system or the VRID of a neighboring VRRP router.
- Backup Routers*
- Backup routers back up the primary router of a specified virtual router and assume Master state responsibilities for the virtual router should the primary router fail.
 - When you define a Backup virtual router, you cannot use the VRID of a primary router that is defined on the system. You cannot define a Primary and Backup VRRP router for the same virtual router on the same routing device.
- Address Mode*
- In `auto-learn` mode, systems learn the IP addresses to associate with the specified VRID.
 - In `IP address` mode, the system prompts you to select the interface index from a list.
 - After a reboot, the address learning process restarts for each virtual router in `auto-learn` address mode.
 - When you define a Primary virtual router, selecting `auto-learn` as the address mode automatically adds all IP addresses that are associated with the selected VLAN to the primary virtual router.

- When you define a Primary router on a VLAN that contains a single interface, the single interface is automatically chosen as the primary address when you select `IP-address` as the Address mode.
 - When you define a Backup virtual router, selecting `auto-learn` as the address mode configures the Backup router to learn the IP addresses that are associated with the virtual router by means of VRRP advertisements from the Primary router. The Primary router must be up for backup routers to auto-learn the addresses that are associated with the specified VRID.
 - When you define Backup virtual routers, the `auto-learn` address mode option enables auto address learning for the specified VRID. If a new interface is added to the VLAN on a primary virtual router, the new IP address is sent out in VRRP advertisements so that the Backup routers in `auto-learn` mode can learn the new address without having to manually add the new address to each backup router.
- Advertisement Intervals*
- The smaller the advertisement interval, the smaller the failover time if the master fails.
 - The advertisement interval must be the same across the set of VRRP routers that are associated with a single VRID. Backup routers must have the same advertisement interval as the Master router.

Options

Prompt	Description	Possible Values	[Default]
Virtual router type	Type of virtual router that you want to define	<ul style="list-style-type: none"> ■ Primary ■ Backup 	Primary
VLAN interface index	Index number of the virtual LAN (VLAN) on which you want to define the virtual router	<ul style="list-style-type: none"> ■ Index number of an IP virtual LAN (VLAN) that is defined on the system. ■ ? (for a list of selectable indexes) 	Index number of first available VLAN
VRID	Virtual router identifier. Identifies the virtual router that you want to define on the LAN.	1 – 255	1
Address mode	Method by which the virtual router you want to define learns its IP addresses	<ul style="list-style-type: none"> ■ auto-learn ■ IP address 	auto-learn

Prompt	Description	Possible Values	[Default]
Advertise interval	Time between virtual router advertisements.	1 – 255 seconds	1
Preempt mode	Whether a higher priority backup router may preempt a lower priority master	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Authentication type	Whether a password is needed to access the virtual router	<ul style="list-style-type: none"> ■ none ■ pass 	none
Password	Character string to authenticate access to virtual router	up to eight alphanumeric characters	–

IP VRRP Define Example

```

Select menu option (ip/vrrp): define
Enter virtual router's type (Primary,Backup) [Primary]:
Enter VLAN interface index {2-5|?}: 2
Enter VRID (1-255) [1]: 2
Enter address mode (auto-learn,IP-address) [auto-learn]:
Enter the advertise interval in sec (1-255) [1]:
Enter virtual router preempt mode (no,yes) [yes]:
Enter Authentication Type (none,pass) [pass]: pass
Enter 8 characters password {?}: echoe

```

ip vrrp modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Modifies an existing virtual router.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`ip v modi`

3900
9300

Important Considerations

- Authentication passwords can be up to eight alphanumeric characters.
 - You can define *one* Primary router per VLAN.
 - Primary routers own the IP addresses that you associate with a virtual router.
 - When you define a Primary virtual router, the possible VLANs that you can select are the IP VLANs on the router that have no virtual routers configured.
 - The virtual router ID (VRID) must be unique across all locally attached LAN segments and unique for the local router.
 - When you define a Primary virtual router, you cannot use the VRID of a virtual router that is already defined on the system or the VRID of a neighboring VRRP router.
 - Backup routers back up the primary router of a specified virtual router and assume Master state responsibilities for the virtual router should the primary router fail.
 - When you define a Backup virtual router, you cannot use the VRID of a primary router that is defined on the system. You cannot define a Primary and Backup VRRP router for the same virtual router on the same routing device.
 - In `auto-learn` mode, systems learn the IP addresses to associate with the specified VRID.
 - In `IP address` mode, the system prompts you to select the interface index from a list.
 - After a reboot, the address learning process restarts for each virtual router in `auto-learn` address mode.
 - When you define a Primary virtual router, selecting `auto-learn` as the address mode automatically adds all IP addresses that are associated with the selected VLAN to the primary virtual router.
- Primary Routers*
- Backup Routers*
- Address Mode*

- When you define a Primary router on a VLAN that contains a single interface, the single interface is automatically chosen as the primary address when you select `IP-address` as the Address mode.
 - When you define a Backup virtual router, selecting `auto-learn` as the address mode configures the Backup router to learn the IP addresses that are associated with the virtual router by means of VRRP advertisements from the Primary router. The Primary router must be up for backup routers to auto-learn the addresses that are associated with the specified VRID.
 - When you define Backup virtual routers, the auto-learn address mode option enables auto address learning for the specified VRID. If a new interface is added to the VLAN on a primary virtual router, the new IP address is sent out in VRRP advertisements so that the Backup routers in auto-learn mode can learn the new address without having to manually add the new address to each backup router.
- Advertisement Intervals*
- The smaller the advertisement interval, the smaller the failover time if the master fails.
 - The advertisement interval must be the same across the set of VRRP routers that are associated with a single VRID. Backup routers must have the same advertisement interval as the Master router.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) on which you want to define the virtual router	<ul style="list-style-type: none"> ■ Index number of an IP virtual LAN (VLAN) that is defined on the system. ■ ? (for a list of selectable indexes) 	Index number of first available VLAN
VRID	Virtual router identifier. Identifies the virtual router that you want to define on the LAN.	1 – 255	1
Virtual router type	Type of virtual router that you want to define	<ul style="list-style-type: none"> ■ Primary ■ Backup 	Primary
Address mode	Method by which the virtual router you want to define learns its IP addresses	<ul style="list-style-type: none"> ■ auto-learn ■ IP address 	auto-learn

Prompt	Description	Possible Values	[Default]
Advertise interval	Time between virtual router advertisements.	1 – 255 seconds	1
Preempt mode	Whether a higher priority backup router may preempt a lower priority master	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Authentication type	Whether a password is needed to access the virtual router	<ul style="list-style-type: none"> ■ none ■ pass 	none
Password	Character string to authenticate access to virtual router	up to eight alphanumeric characters	–

IP VRRP Modify Example

```

Select menu option (ip/vrrp): modify
Enter VLAN interface index {2-3|?}: 2
Enter virtual router ID {1|?} [1]:
Enter virtual router's type (Primary,Backup) [Primary]:
Enter address mode (auto-learn,IP-address) [auto-learn]: IP-address
Old Ip Association address list:
  VRID    VIDX    Address
   1         2    158.101.175.228
Interface 158.101.175.228 will be selected as your primary address.
Enter the advertise interval in sec (1-255) [1]:
Enter virtual router preempt mode (no,yes) [yes]: no
Enter Authentication Type (none,pass): none
Enter virtual router state (enabled,disabled) [enabled]:

```

ip vrrp remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Removes one or more existing virtual routers from the system.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip v r

Important Consideration

- If you attempt to remove a virtual router that is in the Master state, you are prompted to confirm the operation:
 - If you enter `no`, the system does not remove the virtual router.
 - If you enter `yes`, the system removes the virtual router, which sends an advertisement to the other virtual routers that one of them must assume Master responsibilities immediately.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) on which you want to define the virtual router	<ul style="list-style-type: none"> ■ Index number of a IP virtual LAN (VLAN) defined on the system ■ ? (for a list of selectable indexes) 	Index number of first available VLAN
VRID	Virtual router identifier. Identifies the virtual router that you want to define on the LAN	1 – 255	1

IP VRRP Remove Example

```
Select menu option (ip/vrrp): remove
Enter VLAN interface index (2-3|all|?): 2
Enter virtual router ID (1|?) [1]:
```

ip vrrp mode *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Enables or disables a configured virtual router.

Valid Minimum Abbreviation

`ip v mode`

3900
9300

Important Considerations

- You must configure the virtual router before you can enable it.
- You cannot modify or remove a virtual router that is enabled; you must disable the virtual router before you can change or delete the virtual router.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) on which you want to define the virtual router	<ul style="list-style-type: none"> ■ Index number of a IP virtual LAN (VLAN) defined on the system ■ all ■ ? (for a list of selectable indexes) 	Index number of first available VLAN
VRID	Virtual router identifier. Identifies the virtual router that you want to define on the LAN	1 – 255	1
Virtual router mode	Explicitly turns on or turns off a configured virtual router	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled

IP VRRP Mode Example

```
Select menu option: ip vrrp mode
Enter VLAN interface index (2-3|all|?): all
Enter virtual router ID (1-2|all|?): all
Vrid 1 - Enter virtual router mode (enabled,disabled)
[disabled]: enabled
Vrid 2 - Enter virtual router mode (enabled,disabled)
[disabled]: enabled
```


ip vrrp neighbor *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays a list of neighboring virtual routers.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip v n

3900

9300

Important Considerations

- Any locally defined virtual router is not displayed.
- If the Address and MasterRouterAddr fields contain the same IP address, the listed virtual router is in the Master state.

Fields in the IP VRRP Neighbor Display

Field	Description
VLAN Index	Index number of the VLAN on which the virtual router is defined
VRID	Virtual Router ID. Number that identifies the virtual router on the LAN
Address	IP address of the neighbor virtual router, which may be a Master or Backup router
MasterRouterAddr	IP address of the Master virtual router
Interval	Time, in seconds, between virtual router advertisements
Priority	Priority among the backup routers to become the Master virtual router
Auth	Authentication type: whether a password is needed to access the virtual router
Config	Whether the virtual router has been locally configured

ip vrrp statistics *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays general VRRP statistics for the virtual router.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip v st

3900
9300

Fields in the IP VRRP Statistics Display

Field	Description
ckSumErrors	Total number of VRRP advertisements with an invalid VRRP checksum value that this virtual router has received
versionErrors	Total number of VRRP advertisements with an unknown or unsupported version number that this virtual router has received.
vriderrors	Total number of VRRP advertisements with an invalid VRID number that this virtual router has received

IP MULTICAST

This chapter provides guidelines and other key information about how to configure and manage IP multicast routing commands from the Administration Console of the CoreBuilder® 3500 and CoreBuilder 9000 Layer 3 switching modules.



For the CoreBuilder 9000 platform, the commands in this chapter apply to Layer 3 switching modules only.



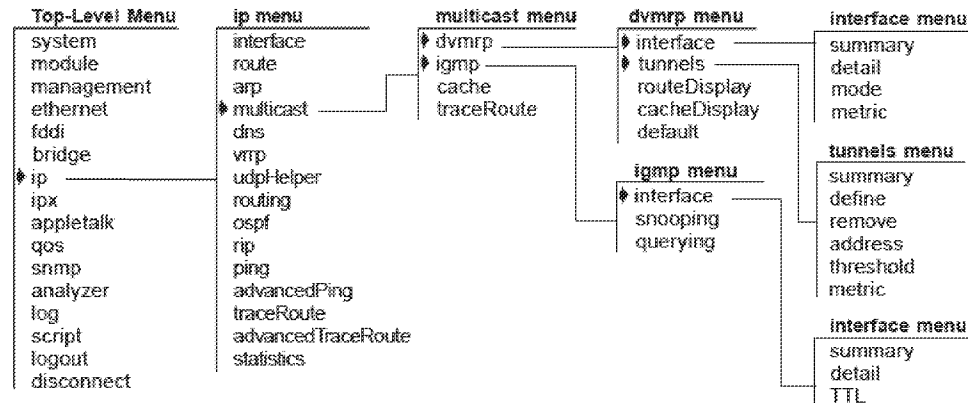
For more information about IP multicast technology, concepts, and implementation procedures, see the Implementation Guide for your system.



For IGMP commands in Layer 2 switching systems (CoreBuilder 9400, CoreBuilder 9000 Layer 2 switching modules, SuperStack® II Switch 3900, and SuperStack II Switch 9300), see Chapter 9.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured on your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



**ip multicast dvmrp
interface summary**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays summary information about IP interfaces that may or may not be operating as IP multicast routing interfaces using the Distance-Vector Multicast Routing Protocol (DVMRP).

Valid Minimum Abbreviation

`ip m d i s`

Fields in the IP Multicast DVMRP Interface Summary Display

Field	Description
Index	Number associated with the interface for identification purposes
Address	IP address of the interface
Metric	Numeric DVMRP metric or "cost" that you assign to the interface
State	Role that the interface plays in IP multicast delivery. One or more of the following descriptors may appear: <ul style="list-style-type: none"> ■ <code>querier</code> — The interface is functioning as the IGMP Querier for its subnetwork. ■ <code>non-querier</code> — The interface is <i>not</i> functioning as the IGMP Querier for its subnetwork. ■ <code>leaf</code> — There are no routers downstream of this interface; IP multicast group members may reside on this subnetwork. ■ <code>non-leaf</code> — The interface is a branch in the IP multicast delivery tree. There are one or more IP multicast routing interfaces downstream of this interface. ■ <code>one-way</code> — Traffic is moving downstream only. ■ <code>disabled</code> — DVMRP is disabled on the interface. ■ <code>up</code> — The IP interface is available to support network communication. ■ <code>down</code> — The IP interface is not available to support network communication.

**ip multicast dvmrp
interface detail****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays information about IP interfaces that run the Distance-Vector Multicast Routing Protocol.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**`ip m d i s`3900
9300**Fields in the IP Multicast DVMRP Interface Detail Display**

Field	Description
Index	Number associated with the interface for identification purposes
Address	IP address of the interface
Metric	Numeric DVMRP metric or “cost” that you assign to the interface
State	Role that the interface plays in IP multicast delivery. One or more of the following descriptors may appear: <ul style="list-style-type: none"> ■ <code>querier</code> — The interface is functioning as the IGMP Querier for its subnetwork. ■ <code>non-querier</code> — The interface is <i>not</i> functioning as the IGMP Querier for its subnetwork. ■ <code>leaf</code> — There are no routers downstream of this interface; IP multicast group members may reside on this subnetwork. ■ <code>non-leaf</code> — The interface is a branch in the IP multicast delivery tree. There are one or more IP multicast routing interfaces downstream of this interface. ■ <code>one-way</code> — Traffic is moving downstream only. ■ <code>disabled</code> — DVMRP is disabled on the interface. ■ <code>up</code> — The IP interface is available to support network communication. ■ <code>down</code> — The IP interface is not available to support network communication.
Group	IP multicast group addresses of the traffic that is being received and forwarded on that interface.
Peer, Port	IP address of the upstream router. The additional information to the right relates to the version of DVMRP that is running and the port in the local interface that connects to the peer router.

**ip multicast dvmrp
interface mode**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables the Distance-Vector Multicast Routing Protocol (DVMRP) per routing interface. This protocol facilitates router-to-router communication for building source-rooted spanning trees that deliver IP multicast traffic to IP multicast group members.

Valid Minimum Abbreviation

ip m d i m

Important Considerations

- When DVMRP is enabled on an interface, the interface is configured with the default value of 1 for the metric, which you can modify at any time. See “ip multicast dvmrp interface metric” later in this chapter.
- If DVMRP is enabled on any interface, IGMP snooping should also be enabled in the system. See “ip multicast igmp snooping” later in this chapter.
- If DVMRP is disabled, the interface cannot participate in building spanning trees for IP multicast. However, as long as IGMP snooping is enabled, the interface forwards appropriate IP multicast traffic to downstream group members. If IGMP snooping is disabled, then the interface only forwards IP multicast traffic with addresses in the reserved range.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the interface for which you want to enable or disable DVMRP	<ul style="list-style-type: none"> ■ A valid IP interface index number ■ all ■ ? (for a list of selectable indexes) 	–
DVMRP mode	Whether DVMRP mode is enabled or disabled	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled (factory default), or current value

**ip multicast dvmrp
interface metric**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the DVMRP metric on an interface for which DVMRP is enabled.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

ip m d i m

Important Considerations

- Use this command if you want to modify the metric value of 1 that the system assigns to an interface when you define it, even if DVMRP is not yet enabled.
- The metric affects the shape of the IP multicast spanning tree when there are multiple paths to the same downstream destination. The lower cost path is the preferred path.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the routing interface for which you want to modify the default metric	<ul style="list-style-type: none"> ■ A valid IP interface index number ■ ? (for a list of selectable index numbers) 	–
metric	DVMRP cost for the interface	<ul style="list-style-type: none"> ■ 1 – 32 	1 (factory default), or current value

**ip multicast dvmrp
tunnels summary**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Summaries key information about DVMRP tunnels that you have configured in your system. Tunnels enable IP multicast spanning trees to be constructed through and beyond areas of the network (routers) that do not support IP multicast routing. The two tunnel end points must lie in different systems and subnetworks.

Valid Minimum Abbreviation

ip m d t s

Important Considerations

- The index number shown in the DVMRP tunnel summary display is the tunnel index number. When you define a DVMRP tunnel, the system assigns a tunnel index number to it, which is different from the routing interface index number. Tunnel index numbers provide a way to identify individual tunnels, which is necessary because multiple tunnel end points can be configured on the same routing interface. Tunnel index numbers are also needed so that you can remove tunnels without removing the interface with which it is associated.
- When you remove a tunnel, the system does not dynamically re-order remaining tunnels in the DVMRP tunnel summary display. For example, if you had three tunnels with tunnel index numbers 1, 2, and 3 and you then removed tunnel 2, the display lists the remaining tunnels with their original tunnel index numbers (1 and 3, in this example). The system assigns tunnel index 2 to the next *new* tunnel that you define. After 2 is used, the system can assign tunnel index 4 for the next new tunnel, and so on.
- You can define multiple IP multicast tunnel end points on the same local routing interface, but each must lead to a different remote interface. You cannot define multiple IP multicast tunnels between the same two end points (interfaces).

Fields in the IP Multicast DVMRP Tunnels Summary Display

Field	Description
Index	Tunnel index number, which is different from the routing interface index number that is shown under <code>Index</code> in other displays.
Local address	IP address of the local interface that serves as one of two multicast tunnel end points.
Remote address	IP address of the remote interface (a different system, a different subnetwork) that serves as the other multicast tunnel end point.
Metric	DVMRP cost of the tunnel. The system assigns a value of 1 when you define the tunnel, but you can modify that value at any time (see “ip multicast dvmrp tunnels metric”). This value can be different from the metric that you assigned to the interface itself (see “ip multicast dvmrp interface metric”).
TTL	Time-to-live (TTL) threshold of the tunnel. The system assigns a value of 1 when you define the tunnel, but you can modify that value at any time (see “ip multicast dvmrp tunnels threshold”). This value can be different from the TTL threshold that you assigned to the interface itself (see “ip multicast igmp interface TTL”).
State	Role that the interface in the multicast delivery tree. For possible entries and definitions, see “ip multicast dvmrp interface summary” earlier in this chapter.

**ip multicast dvmrp
tunnels define**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines one end point of a DVMRP tunnel. The other tunnel end point lies on an IP multicast routing interface on a different system and subnetwork. One or more unicast routers lie between these tunnel end points.

Valid Minimum Abbreviation

`ip m d t d`

Important Considerations

- IP multicast tunnels are not required in all networks. Configure a tunnel only if you need to have IP multicast traffic forwarded through one or more routers that do not understand IP multicast protocols and would therefore filter IP multicast packets. Because IP multicast packets are encapsulated in unicast format at the tunnel entrance point, the interim routers in the tunnel forward the packets onward toward the other tunnel exit point.
- Think of an IP multicast tunnel end point as being layered on top of a regular DVMRP routing interface. Therefore, before you can define a multicast tunnel end point in your system, you must first define at least one IP virtual LAN (VLAN), define at least one IP interface, and enable DVMRP on the interface.
- The remote tunnel end point must lie on a different system and subnetwork.
- You must define the tunnel on both end points — that is, on both the local system and the remote system — even though you specify the address of the remote interface in the local system.
- When you define a tunnel with local and remote addresses, the system automatically assigns the value 1 as both the tunnel metric and the tunnel TTL threshold, as shown in the IP multicast DVMRP tunnel summary display. You can change these values through menu options.
- IP multicast interfaces and tunnels have similar characteristics, such as TTL threshold and metric. The characteristics of a tunnel do not have to match the characteristics of the interface on which it is configured.
- You can define multiple tunnel end points on the same local routing interface in your system, but these tunnels must lead to different remote routing interfaces.

Options

Prompt	Description	Possible Values	[Default]
interface	Index number of the interface on which you want to create a DVMRP tunnel end point	<ul style="list-style-type: none">■ A valid IP interface index number■ ? (for a list of selectable indexes)	–
Remote address	IP address of the remote multicast tunnel end point. Use standard dotted decimal notation.	A valid IP interface on a different system and subnetwork	–

**ip multicast dvmrp
tunnels remove****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Deletes a DVMRP tunnel end point from the system.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**

ip m d t r

Important Considerations

- To remove a tunnel, specify its tunnel index number. This number is different from the routing interface index number. Reference the DVMRP tunnel summary display prior to deleting a tunnel.
- If you try to remove an IP interface in your system, and you have a DVMRP tunnel defined on that interface, the system warns you with an error message. Before you can remove the IP interface, you must remove the DVMRP tunnel.
- When you remove a tunnel, the system does not dynamically re-order remaining tunnels in the DVMRP tunnel summary display. For example, if you had three tunnels with tunnel index numbers 1, 2, and 3 and you then removed tunnel 2, the display lists the remaining tunnels with their original tunnel index numbers (1 and 3, in this example). The system assigns tunnel index 2 to the next *new* tunnel that you define. After 2 is used, the system can assign tunnel index 4 for the next new tunnel, and so on.

Options

Prompt	Description	Possible Values	[Default]
Multicast tunnel index	Index number of the multicast tunnel that you want to remove from the system	<ul style="list-style-type: none"> ■ A valid DVMRP tunnel index number ■ ? (for a list of selectable tunnel index numbers) 	–

**ip multicast dvmrp
tunnels address**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the remote IP address that is defined in an existing DVMRP tunnel.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ip m d t a`

Important Consideration

- The remote address that you specify must represent a routing interface on a different system and subnetwork.

Options

Prompt	Description	Possible Values	[Default]
tunnel	Index number of the tunnel for which you modify the remote tunnel end point	<ul style="list-style-type: none"> ■ A valid DVMRP tunnel index number in the system ■ ? (for a list of selectable tunnel index numbers) 	–
remote address	A valid IP address on a different system and subnetwork. Use the 0.0.0.0 format.	A valid IP address	current value

ip multicast dvmrp tunnels threshold

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the time-to-live (TTL) threshold on an existing DVMRP tunnel.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Configuration

`ip m d t t`

Important Consideration

- When you first define a tunnel, the system automatically assigns the value 1 as the TTL threshold for the tunnel (which is different from the interface TTL threshold). Use this command to modify the TTL threshold value on any existing tunnel.

- 3900
- 9300

Options

Prompt	Definition	Possible Values	[Default]
tunnel	Index number of the existing DVMRP tunnel on which you want to modify the TTL threshold	<ul style="list-style-type: none"> ■ A valid DVMRP tunnel index number ■ ? (for a list of selectable tunnel index numbers) 	–
threshold	Value that determines whether IP multicast packets are forwarded. The interface compares the packet TTL to the TTL threshold	1 – 32	1 (factory default), or current value

**ip multicast dvmrp
tunnels metric**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the metric or “cost” of an existing DVMRP tunnel.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Configuration

`ip m d t m`

Important Consideration

- When you first define a tunnel, the system automatically assigns the value 1 as the metric or “cost” of the tunnel (which is different from the interface metric). Use this command to modify the metric value on any existing tunnel.

Options

Prompt	Definition	Possible Values	[Default]
tunnel	Index number of the existing DVMRP tunnel on which you want to modify the metric	<ul style="list-style-type: none"> ■ A valid DVMRP tunnel index number ■ ? (for a list of selectable tunnel index numbers) 	–
metric	DVMRP cost for the tunnel. This value affects the shape of the IP multicast spanning tree when there are multiple paths to the same downstream destination. The lower cost path is chosen first.	1 – 32	1 (factory default), or current value

**ip multicast dvmrp
routeDisplay**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays IP multicast route information that your system has learned from using the Distance-Vector Multicast Routing Protocol (DVMRP). The system uses this information to forward IP multicast traffic that it receives.

Valid Minimum Abbreviation

ip m d r

Fields in the IP Multicast DVMRP Route Display

Field	Description
Origin	IP address of the subnetwork that contains an IP multicast source, followed by a forward slash and subnetwork mask.
Gateway	IP address of the routing interface that lies upstream of the local system on the path back towards an IP multicast source. If the source subnetwork is connected directly to your system, this field contains a dash (--).
Metric	Number of hops from your system back to the origin subnetwork. This value is <i>not</i> the DVMRP interface or tunnel metric, which are shown under <code>Metric</code> in other displays. Occasionally, instead of a numeric value, you may see <code>NR</code> , meaning "network unreachable." Your system may have trouble computing the hop count because of factors such as an upstream router being temporarily congested. This condition is usually resolved in a short period of time.
Tmr	Amount of time (in seconds) since each entry was last reset.
Parent	The interface that connects to the upstream router (Gateway). Because DVMRP forms a loopless spanning tree to reach all hosts for a given IP multicast group, your system always chooses a single parent interface. Either an <code>I</code> or a <code>T</code> precedes the index number. An <code>I</code> indicates that the index is an interface index number. A <code>T</code> indicates that the index is a tunnel index number.
Children	Interfaces that communicate with downstream routers or local subnetworks. The system forwards incoming IP multicast traffic through these interfaces. Either an <code>I</code> or a <code>T</code> precedes each index number. An <code>I</code> precedes an interface index number. A <code>T</code> precedes a tunnel index number.

**ip multicast dvmrp
cacheDisplay**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the DVMRP cache, which is a collection of information about the IP multicast packets that have traveled through the system.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

`ip m d c`

Options

Prompt	Description	Possible Values	[Default]
Multicast source address	Source for which you want to view cache information	<ul style="list-style-type: none"> Depends on your network 255.255.255.255 for all sources 	255.255.255.255 (factory default), or current value
Multicast group address	Multicast group for which you want to view cache information	<ul style="list-style-type: none"> Depends on your network 255.255.255.255 for all groups 	255.255.255.255 (factory default), or current value

Fields in the IP Multicast DVMRP Cache Display

Field	Description
Source	Information about IP multicast sources: <ul style="list-style-type: none"> Entries preceded by angle brackets (>) are subnetworks that contain sources. Entries without angle brackets are the IP addresses of source devices.
Group	IP multicast group address of packets coming from the source and subnetwork to the left.
CTmr	Time since the cache entry was originally recorded. Time is noted in hours (h), minutes (m), and seconds (s).
Age	Value that indicates the remaining life for the cache entry. Time is recorded in minutes (m) and seconds (s). The system assigns a life of approximately 7 minutes to each entry. When the age of the entry decreases to zero, the entry either disappears or is refreshed.
PTmr	Time remaining before the system sends a prune message to an upstream router. Time is shown in minutes (m) and seconds (s). When traffic is actively flowing, a dash (-) indicates that no prune message has been sent upstream.

Field	Description
inVif	<p>Interface that receives incoming IP multicast traffic from the spanning tree for the source, subnetwork, and group listed on the left.</p> <p>The interface is presented as an index number and either an I or a T precedes the index number. An I precedes a routing interface index number. A T precedes a tunnel index number.</p> <p>A P after the index number indicates that a prune message has been sent to an upstream router.</p> <p>The entry <code><none></code> may appear if the system is not able to build the cache entry correctly. This temporary condition corrects itself quickly.</p>
outVif	<p>Interfaces to which traffic from the inVif is being forwarded.</p> <p>Each interface is presented as an index number and either an I or a T precedes each index number. An I precedes a routing interface index number. A T precedes a tunnel index number.</p> <p>A p after an index number indicates that the upstream router has pruned this branch of the delivery tree and no multicast packets are being forwarded through this local interface. Eventually this entry disappears from the cache display.</p> <p>Either no entry or <code><none></code> appears in this column if the system is not able to build the cache entry correctly. This temporary condition corrects itself quickly.</p>
Ports	<p>Physical ports that correspond to the interfaces that are listed in the outVifs field. The Ports field shows a dash (--) when there are no outgoing interfaces and when the outgoing interfaces are tunnels.</p>

**ip multicast dvmrp
default**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Configures a default route for IP multicast traffic on a DVMRP interface. This interface advertises itself as a default route to neighboring DVMRP routers.

Valid Minimum Abbreviation

`ip m d d`

Important Considerations

- A default route metric of 0 means that the default route function is not activated on the interface (interface does not advertise 0.0.0.0 to DVMRP routers). Values other than 0 means that the default route function is activated and these values represent the “cost” of the default route.
- Definitions of default route modes:
 - **all** — The interface advertises the default route plus all other known routes to neighboring DVMRP routers.
 - **only** — The interface advertises only the default route to neighboring DVMRP routers.

If the system learns a default route, it propagates it no matter which mode is set on a given interface.
- The system allows you to configure an interface as a DVMRP default route, even when DVMRP is disabled on the interface. If DVMRP is disabled, the interface does not advertise itself as a default route.

Options

Prompt	Definition	Possible Values	[Default]
interface	Index number of the routing interface on which you want to configure a default route	<ul style="list-style-type: none"> ■ A valid interface index number ■ ? (for a list of selectable indexes) 	1 (factory default), or current value
default route metric	Value that you assign to the default route as the “cost” of that route	0 – 32	0 (factory default), or current value
default route advertise mode	Routes that the interface advertises to neighboring DVMRP routers	<ul style="list-style-type: none"> ■ all ■ only 	all (factory default), or current value

**ip multicast igmp
interface summary*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Summarizes key information about IGMP interfaces.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**`ip m i i s`3900
9300**Fields in the IP Multicast IGMP Interface Summary Display**

Field	Description
Index	Number assigned to the routing interface to its right.
Address	IP address of a routing interface in the system
TtlThreshold	Time-to-live (TTL) threshold that is assigned to the interface. This threshold affects IP multicast packets only.
Protocol	Multicast routing protocol that registers with IGMP. In release 3.0 software, there is one supported routing protocol (DVMRP).
Querier	IP address of the IGMP querier in the subnetwork to which the interface belongs. If the interface is functioning as the IGMP querier, this field shows <code>Self</code> .

**ip multicast igmp
interface detail****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Supplements the IP multicast IGMP interface summary display with group and port information.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**`ip m i i d`**Fields in the IP Multicast IGMP Interface Detail Display**

Field	Description
Index	Number assigned to the routing interface to its right for identification purposes.
Address	IP address of a routing interface in the system that is identified by the index number to its left.
TtlThreshold	Time-to-live (TTL) threshold that is assigned to the interface. This threshold affects IP multicast packets only.
Protocol	Multicast routing protocol that registers with IGMP. In release 3.0 software, there is one supported routing protocol (DVMRP).
Querier	IP address of the IGMP querier in the subnetwork to which the interface belongs. If the interface is functioning as the IGMP querier, this field shows <code>self</code> .
group	IP multicast group address for which packets have been received or forwarded
port(s)	Physical port numbers that are associated with the interface listed in the <code>Address</code> field that see incoming or outgoing traffic.

**ip multicast igmp
interface TTL**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the time-to-live (TTL) threshold of a given routing interface. The interface compares the TTL value in each IP multicast packet against its TTL threshold. If the packet TTL is greater than the threshold TTL, the interface decrements the packet TTL by 1 and forwards the packet, provided that no other restrictions exist.

Valid Minimum Abbreviation

```
ip m i i t
```

Important Considerations

- Because IGMP is enabled by factory default, the system assigns a TTL threshold value of 1 as soon as you create an IP interface.
- This TTL threshold affects IP multicast packets only.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index numbers of the interfaces for which you want to modify the TTL threshold	<ul style="list-style-type: none"> ■ One or more valid interface index numbers ■ ? (for a list of selectable indexes) 	–
TTL threshold	Value you want to assign to the specified interfaces	0 – 255	1 (factory default), or current value

**ip multicast igmp
snooping*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Enables or disables the system's ability to understand the Internet Group Management Protocol (IGMP) and snoop on IGMP packets to determine if IP multicast group members exist downstream from routing interfaces and therefore if the system should forward group traffic on those interfaces.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation`ip m i s`**Important Considerations**

- Your selection applies to all interfaces in the system.
- 3Com recommends that you keep IGMP snooping enabled at all times. It adds little processing overhead to the system and enhances the efficiency of your network if IP multicast traffic is present.

Options

Prompt	Description	Possible Values	[Default]
snooping mode	Whether the system can observe, record, and react to IGMP packets and set filters on appropriate ports in an interface	<ul style="list-style-type: none"> ■ enabled ■ disabled 	enabled (factory default), or current value

**ip multicast igmp
querying**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables the system's ability to operate as the Internet Group Management Protocol (IGMP) querier if so elected by other IGMP-capable devices in the subnetwork. The IGMP querier is always the device with the lowest IP address.

Valid Minimum Abbreviation

`ip m i q`

Important Considerations

- Your selection applies to all interfaces in the system.
- The most efficient bandwidth usage is achieved by having the device that is closest to the source of IP multicast traffic operate as the querier for a given subnetwork.

Options

Prompt	Description	Possible Values	[Default]
query mode	Whether the system can offer itself as a candidate for election as the IGMP querier	<ul style="list-style-type: none"> ■ enabled ■ disabled 	enabled (factory default), or current value

ip multicast cache

✓ 3500

✓ 9000

9400

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays information about IP multicast traffic that has been observed on the system. For more detailed information, review the DVMRP cache. (See “ip multicast dvmrp cacheDisplay” earlier in this chapter.)

Valid Minimum Abbreviation

ip m c

3900

9300

Important Consideration

- Although the Administration Console menu description is `protocol independent multicast cache`, this cache is not related to the multicast routing protocol called *Protocol Independent Multicast (PIM)*.

Options

Prompt	Description	Possible Values	[Default]
Multicast source address	Source for which you want to view cache information	<ul style="list-style-type: none"> ■ Depends on your network ■ 255.255.255.255 for all sources 	255.255.255.255 (factory default), or current value
Multicast group address	Multicast group for which you want to view cache information	<ul style="list-style-type: none"> ■ Depends on your network ■ 255.255.255.255 for all groups 	255.255.255.255 (factory default), or current value

Fields in the IP Multicast Cache Display

Field	Description
source	Subnetwork that contains a source device that is sending traffic addressed to the IP multicast group listed in the <code>group</code> field.
group	IP multicast group address of packets coming from the subnetwork listed to its left.
inVif	Index number of the interface that receives incoming IP multicast group traffic. Either an <code>I</code> or a <code>T</code> precedes the index number. An <code>I</code> indicates a regular IP multicast interface. A <code>T</code> indicates that the interface also operates as a DVMRP tunnel.
outVif	Index numbers of the interfaces to which traffic from the <code>inVif</code> is being forwarded.
inPorts	Physical port that corresponds to the interface that is listed in the <code>inVifs</code> field.
outPorts	Physical ports that correspond to the interfaces that are listed in the <code>outVifs</code> field.

**ip multicast
traceRoute****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Provides a method for tracing the path that an IP multicast packet takes from a source to a particular receiver. Unlike unicast IP traceroute, multicast traceroute works in the reverse and requires a special packet type and implementation in routing devices.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation`ip m t`**Important Considerations**

- This command traces the path backwards from a specific receiving device to a specific source device. When you use this command, the receiver is assumed to be the system to which you are connected.
- This command produces a display that shows IP addresses of the interfaces that span from your system back to the source that you specify. The display also shows the number of hops back to those interfaces, the multicast routing protocols used, and the amount of time it takes to reach each hop from the receiver.
- All interim devices must support IP multicast traceroute for you to see a complete path on the display.

Options

Prompt	Description	Possible Values	[Default]
source IP address	IP address of the source device that sends traffic to a specific IP multicast group address	Any valid IP address for IP multicast source devices in your network	–
multicast group address	The IP multicast group address that the source is using for a particular application. This is useful when all applications come from the same source.	Any valid IP multicast group address used by source devices in your network	–

19

OPEN SHORTEST PATH FIRST (OSPF)

This chapter describes commands that you can use to configure Open Shortest Path First (OSPF) routing on your system.



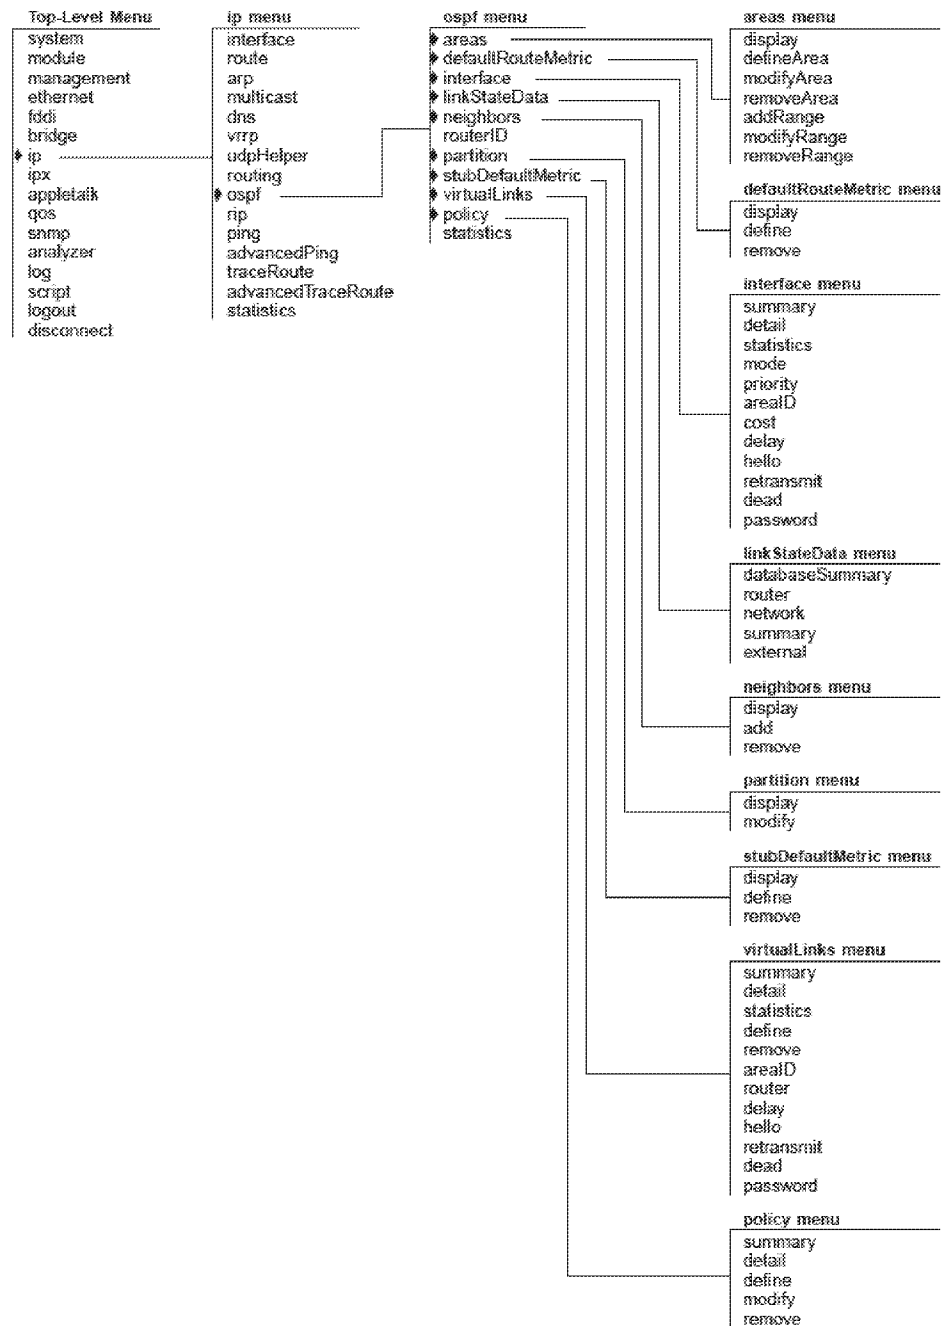
For more information about administering OSPF routing on your network, see the Implementation Guide for your system.



For the CoreBuilder® 9000, the commands in this chapter apply to Layer 3 switching modules only.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



ip ospf areas display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays a list of existing OSPF areas.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip o a di

3900

9300

Fields in the IP OSPF Areas Display

Field	Description
Advertise	Whether the network range is advertised (y) or not (n)
AreaID	Area identifier
Indx	Entry index number for the area
IP Address	Network portion of IP address range
Mask	IP address range subnet mask
Stub	Whether the area is a stub area (y) or not (n)

**ip ospf areas
defineArea**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines an OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a de

Important Considerations

- The backbone area 0.0.0.0 is configured by default.
- The area ID must be unique for the autonomous system.
- On the CoreBuilder 3500, you can define a maximum of eight areas.

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Area ID	In the form n.n.n.n (where 0 <= n <= 255); functions as an area identification number to the OSPF autonomous system	Up to 255.255.255.255	–
Stub area	Whether this area is a stub area	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n (factory default), or current value

ip ospf areas
modifyArea

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies an existing OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a modifya

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Area	Index number of the area that you want to modify	<ul style="list-style-type: none"> ■ Valid area index number ■ ? (for a list of selectable indexes) 	–
Area ID	In the form n.n.n.n (where $0 \leq n \leq 255$); functions as an area identification number to the OSPF autonomous system	Up to 255.255.255.255	–
Stub area	Whether this area is a stub area	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n (factory default), or current value

**ip ospf areas
removeArea**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes an existing OSPF area.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip o a removea

Options

3900

9300

Prompt	Description	Possible Values	[Default]
Area	Index number of the area that you want to remove	<ul style="list-style-type: none"> ■ Valid area index number ■ all ■ ? (for a list of selectable indexes) 	First available index number

ip ospf areas
addRange

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Adds a range to an existing OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a a

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Area	Index number of the area to which you want to add the range	<ul style="list-style-type: none"> ■ Valid area index number ■ ? (for a list of selectable indexes) 	–
IP address	IP address of the range that you want to add to the area	Up to 255.255.255.255	–
Subnet mask	Subnet mask of the range that you want to add to the area	Variable, based on address range class	Variable, based on address range class
Advertise range	Whether to advertise area range	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y

**ip ospf areas
modifyRange**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies an OSPF area range.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a modifyr

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Area	Index number of the area that contains the range that you want to modify	<ul style="list-style-type: none"> ■ Valid area index number ■ ? (for a list of selectable indexes) 	–
IP address of range	Existing range that you want to modify (in the form of an IP address)	Up to 255.255.255.255	–
IP address	Range (in the form of an IP address)	Up to 255.255.255.255	Current value
Subnet mask	Subnet mask of the range that you want to modify	Variable, based on address range class	Current value
Advertise range	Whether to advertise the area range	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	Current value

IP OSPF Areas Modify Range Example

```
Select area {1-2|?}: 1
Enter IP address of range to modify: 3.3.3.1
Enter IP address [3.3.3.1]: 2.2.2.2
Enter subnet mask [255.0.0.0]: 255.255.0.0
Advertise this area range (yes,no) [yes]: y
```

ip ospf areas
removeRange

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes an OSPF area range.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a remove

Options

Prompt	Description	Possible Values	[Default]
Area	Index number of the area that contains the range that you want to delete	<ul style="list-style-type: none"> ■ Valid area index number ■ ? (for a list of selectable indexes) 	–
IP address	IP address of the range that you want to delete	Up to 255.255.255.255	–

**ip ospf
defaultRouteMetric
display**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the cost of a default route.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o d di

Important Considerations

- If a default metric is not defined, the router does not advertise itself as the default router.
- By default, the default route metric is not defined.

Field in the IP OSPF Default Route Metric Display

Field	Description
Default route metric	Cost (metric) that is associated with the default route. A higher cost indicates a slower route, for example, because it entails more hops or less bandwidth.

ip ospf
defaultRouteMetric
define

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines the default route metric for the router.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o d de

Important Considerations

- If a default metric is not defined, the router does not advertise itself as the default router.
- By default, the default route metric is not defined.
- Defining is default route metric is useful when the configuration supports multiple paths to the same destination. It provides a way to signify which of the paths is to be preferred.

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Default route metric	Cost (metric) that is associated with the default route	1 – 65535	–


```
ip ospf
defaultRouteMetric
remove
```

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes the default route metric.

Valid Minimum Abbreviation

```
ip o d r
```

✓ 3500

✓ 9000

9400

3900

9300

Important Considerations

- If a default metric is not defined, the router does not advertise itself as the default router.
- By default, the default route metric is not defined.
- The default route metric is removed immediately after you enter the command. You are not prompted to confirm the deletion.

**ip ospf interface
summary*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays summary information for the system's OSPF interface configuration.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**

ip o i su

3900
9300**Fields in the IP OSPF Interface Summary Display**

Field	Description
ArealD	OSPF area to which the interface belongs
Dead Intvl	Time interval (in seconds) before OSPF declares that a neighbor is dead
Hello Intvl	OSPF Hello packet transmit interval (in seconds) for the interface
Indx	Interface entry index; same number as the IP interface index
Password	Password that is associated with the OSPF interface
Pri	OSPF router priority for the interface
Rxmit Intvl	LSA retransmit interval (in seconds)
Xmit Cost	Interface transmit cost
Xmit Delay	Interface transmit delay (in seconds)

**ip ospf interface
detail****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays summary and detailed information for the system's OSPF interface configuration.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**`ip o i det`3900
9300**Important Consideration**

- The display also indicates whether IP routing and Internet Control Message Protocol (ICMP) router discovery are enabled and gives the OSPF router ID.

Fields in the IP OSPF Interface Detail Display

Field	Description
AreaID	OSPF area to which the interface belongs
BDR	IP interface of the backup designated router (BDR)
Dead Intvl	Time interval (in seconds) before OSPF declares that a neighbor is dead
DR	IP interface of the designated router (DR)
Hello Intvl	OSPF Hello packet transmit interval (in seconds) for the interface
Indx	Index number that corresponds to the IP interface for which OSPF information is displayed
IP address	IP address of the OSPF interface
Notes	When RouterID appears, the interface address is being used as the OSPF router ID
Password	Password that is associated with the OSPF interface
Pri	OSPF router priority for the interface
Rxmit Intvl	LSA retransmit interval (in seconds)

Field	Description
State	<p>Interface state:</p> <ul style="list-style-type: none"> ■ Disabled — OSPF is not enabled on the interface. ■ Down — Interface is down, but OSPF is enabled on it. ■ Loopback — Interface is a loopback interface. ■ Waiting — Router is trying to determine the identity of the DR and BDR on the network. ■ PTP — Interface is operational and connects to either a point-to-point network or a virtual link. The router attempts to form adjacency with the neighboring router. ■ DRother — Interface is on a multiaccess network where this router is not the designated router or backup designated router. ■ BDR — Router is the backup designated router on the attached network. ■ DR — Router is the designated router on the attached network.
Xmit Cost	Interface transmit cost
Xmit Delay	Interface transmit delay (in seconds)

ip ospf interface statistics**For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays statistics that are associated with specified OSPF interfaces.

✓ 3500
 ✓ 9000
 9400

Valid Minimum Abbreviation

ip o i st

Options

3900
 9300

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the interface for which you want to display statistics	<ul style="list-style-type: none"> ■ Valid interface index number ■ all ■ ? (for a list of selectable indexes) 	–

Fields in the IP OSPF Interface Statistics Display

Field	Description
adjacencyDown	Number of times that OSPF adjacencies have gone down
adjacencyUp	Number of times that OSPF adjacencies have been formed
authError	Number of packets discarded due to OSPF authentication errors Interpretation: <ul style="list-style-type: none"> ■ A non-zero value is bad and means that packets from some OSPF routers are being discarded due to authentication errors. This statistic is incremented under the following circumstances: <ul style="list-style-type: none"> ■ If the OSPF packet authentication type is something other than simple password (i.e., cryptographic authentication is not supported in the current implementation). ■ If the OSPF packet contains a password but the interface does not have a password configured. ■ If the OSPF packet has a simple password that does not match the password defined for the OSPF interface.
computedDR	Number of times that the designated router has been computed
lsaXsumError	Number of LSA checksum errors that were detected
mismatchAreaID	Number of interface area ID mismatches that were detected
mismatchAreaType	Number of interface area type mismatches that were detected

Field	Description
mismatchDead	<p>Number of router dead interval mismatches that were detected</p> <p>Interpretation:</p> <ul style="list-style-type: none"> A non-zero value is bad and means that some OSPF routers on the interface are configured with a different dead interval than this router. This prevents the router from becoming a neighbor with these other routers. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> When an OSPF Hello packet is received and the dead interval it defines is different from the dead interval configured on the OSPF interface.
mismatchHello	Number of Hello packet interval mismatches that were detected
mismatchMask	Number of subnet mask mismatches that were detected
packetXsumError	Number of packet checksum errors since interface has come up
receiveDD	<p>Number of database description packets that were received from valid OSPF neighbors.</p> <p>Interpretation:</p> <ul style="list-style-type: none"> A non-zero value is OK. <p>Database description packets are sent when forming adjacencies with valid neighbors. A large number of receiveDD packets in a network whose configuration has not changed could indicate that adjacencies are being torn down and re-established.</p> <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> When an OSPF database descriptor packet from a valid OSPF neighbor is received.
receivedUnknown	Number of unknown LSAs that were received

Field	Description
receiveError	<p>Number of general receive errors.</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value indicates that OSPF packets are being dropped and that this could be causing routing problems. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello packet is received and the packet length is too short. ■ When an OSPF Hello packet is received that has the same router ID as the router receiving the packet. ■ When an OSPF database descriptor packet is received and the packet length is too short. ■ When an OSPF link state request (LSR) packet is received and the packet length is too short. ■ When processing an LSR packet, if the area is not configured on the interface. ■ When an OSPF link state update (LSU) packet is received and the packet length is too short. ■ When processing an LSU packet, if there are more than 500 advertisements the packet is not processed. ■ When an OSPF link state acknowledgement (LSAck) packet is received and the packet length is too short. ■ When processing an LSAck packet, if the area described by the packet is not known by the router receiving the packet. ■ When processing any OSPF packet, if the packet length is less than the OSPF header length then it must have been truncated and the packet is dropped. ■ When an OSPF packet is received on an interface that is not running OSPF. ■ When an OSPF packet is received over a virtual link, but the virtual link is down or not configured. ■ When an OSPF packet is received (over a non-virtual link) from a source whose IP network does not match the IP network of the interface on which it was received. ■ When an OSPF packet is received on a Non-Broadcast Multiple Access network from an unknown neighbor. ■ When an OSPF packet is received whose version is not OSPF version 2.
receiveHello	Number of Hello packets that were received
receiveLsAck	Number of LSA acknowledgments that were received
receiveLSR	Number of LSA request packets that were received

Field	Description
receiveLSU	Number of link state update packets that were received
transmitDD	<p>Number of database description packets that were transmitted</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value is OK. <p>Database description packets are sent when forming adjacencies with valid neighbors. A large number in a network whose configuration has not changed could indicate that adjacencies are being torn down and re-established.</p> <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF database descriptor packet is transmitted.
transmitError	<p>Number of general transmit errors</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value indicates that an OSPF packet could not be sent either out a particular interface, or to a particular destination. This could prevent OSPF from running properly within the autonomous system and lead to routing problems. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello, LSU, or LSAck is being sent as a multicast packet on a non-broadcast multiple access network.
transmitHello	Number of Hello packets that were transmitted
transmitLSAck	Number of LSA acknowledgments that were transmitted
transmitLSR	Number of LSA request packets that were transmitted
transmitLSU	Number of link state update packets that were transmitted

ip ospf interface mode *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*
Enables or disables OSPF on specified IP interfaces.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o i m

Options

3900
9300

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more IP interfaces on which you want to enable or disable OSPF	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
OSPF mode	Whether to disable or enable OSPF on the specified IP interface	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled (factory default), or current value

ip ospf interface priority *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*
Assigns interface priority to the OSPF router.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o i pr

Important Consideration

- The interface priority of an OSPF router determines its status as a designated router.

3900
9300

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more IP interfaces to which you want to assign a priority	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Priority	Interface priority: <ul style="list-style-type: none"> ■ If 0, router will not be the default router. ■ If 1 – 255, the highest priority becomes the designated router. 	<ul style="list-style-type: none"> ■ 0 – 255 	1

**ip ospf interface
areaID**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Associates an interface with an OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o i a

Important Considerations

- Set the area ID to the same value for all routers on the network segment because they are in the same area.
- 0.0.0.0 indicates the OSPF backbone area.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces that you want to associate with the area	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Area ID	ID of area, in the form n.n.n.n (where 0 <= n <= 255) with which you want to associate the specified interfaces	Valid area ID	0.0.0.0 (factory default), or current value

ip ospf interface cost *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Assigns a cost to an OSPF interface.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o i c

Important Consideration

- The interface cost reflects the line speed of the port. Although the system calculates a default cost value based on the module media type, you can use this command to manually change the cost to a different value.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces to which you want to assign a cost	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Cost	Cost that you want to assign to the specified interface (Higher values are slower ports.)	1 – 65535	Cost of slowest port (usually 1)

**ip ospf interface
delay****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Sets the OSPF interface transmit delay.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**`ip o i del`**Important Considerations**

- The system adds the value of the transmit delay to all link state advertisements (LSAs) that it sends out to the network. Set the transmit delay according to the link speed: use a longer transmit delay time for slower link speeds.
- The transmit delay must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces for which you want to set the transmit delay	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Transmit delay	Delay (in seconds) that you want to assign to the specified interface	1 – 65535 seconds	1 (factory default), or current value

ip ospf interface hello *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Sets the interface Hello interval.

Valid Minimum Abbreviation

ip o i he

Important Considerations

- Hello packets inform other routers that the sending router is still active on the network.
- If a router does not send Hello packets for a period of time specified by the dead interval, the router is considered inactive by its neighbors.
- The Hello packet interval must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces for which you want to set the Hello interval	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Hello packet interval	Interval (in seconds) at which the interface transmits Hello packets	1 – 65535 seconds	10 (factory default), or current value

**ip ospf interface
retransmit**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Specifies the OSPF link state advertisement (LSA) retransmit interval for an interface.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o i r

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces on which you want to set the LSA retransmit interval	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
LSA retransmit time	Interval at which the specified interface retransmits LSAs	1 – 65535 seconds	5 (factory default), or current value

ip ospf interface dead *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Specifies the dead interval for an interface.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip o i dea

3900

9300

Important Consideration

- Set the dead interval to the same value for all routers on the network.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces on which you want to set the dead interval	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Dead interval	Maximum duration (in seconds) that neighbor routers wait for a Hello packet before they determine that the transmitting router is inactive	1 – 65535 seconds	40 (factory default), or current value

**ip ospf interface
password**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets password security for an OSPF interface.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o i pa

Important Considerations

- To remove a previously assigned password, set the password to `none`.
- The password must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces for which you want to assign or remove a password	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Password	Password for the specified interface The none option removes a previously assigned password.	<ul style="list-style-type: none"> ■ Up to eight ASCII characters 	none (factory default), or current value

**ip ospf linkStateData
databaseSummary****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Summarizes link state advertisements (LSAs) in the link state database.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**

ip o l d

Important Consideration

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
Area ID	Area ID (in the form n.n.n.n where 0 ≤ n ≤ 255) that corresponds to the OSPF area for which you want to view LSA summary information	Valid area ID	0.0.0.0 (factory default), or current value
Area mask	Subnet mask of OSPF area for which you want to view LSA summary information	Valid area mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data Database Summary Display

Field	Description
Checksum summation	Total of all LSA checksums
External LSAs	Number of external link LSAs
LSA count	Number of LSAs
Network LSAs	Number of network link LSAs
Router LSAs	Number of router link LSAs
Summary LSAs	Number of summary link LSAs

**ip ospf linkStateData
router**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays router link state advertisements (LSAs) in the link state database.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o l r

Important Consideration

- To view link state database information, OSPF must be active (enabled).

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Area ID	Area ID (in the form n.n.n.n where 0 <= n <= 255) that corresponds to the OSPF area for which you want to view router link state advertisement information	Valid area ID	0.0.0.0 (factory default), or current value
Area mask	Subnet mask of OSPF area for which you want to view router link state advertisement information	Valid area mask	0.0.0.0 (factory default), or current value
LSID	Link State ID: router ID of the originating router (in the form of an IP address)	Router ID	0.0.0.0 (factory default), or current value
LSID mask	Link State ID bit mask (Example: 255.0.0.0)	Link State ID bit mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data Router Display

Field	Description
Flags	<ul style="list-style-type: none"> ■ V — Router is the endpoint of an active virtual link that is using the area as a transmit area. ■ ASBR — Router is an autonomous system boundary router. ■ ABR — Router is an area border router.

Field	Description
Link Data	<ul style="list-style-type: none"> ■ PTP — MIB II index value for an unnumbered point-to-point interface. ■ Transit Net — IP address of the router's interface ■ Stub Net — Network IP address mask ■ Virtual link — IP interface address of neighboring router
Link ID	<ul style="list-style-type: none"> ■ PTP — Neighboring router's router ID ■ Transit Net — Address of designated router ■ Stub Net — IP network/subnetwork number ■ Virtual link — Neighboring router's router ID
Link Type	<ul style="list-style-type: none"> ■ PTP — Connection is point-to-point to another router. ■ Transit Net — Connection is to a transit network (one that has more than one OSPF router on it). ■ Stub Net — Connection is to a stub network. ■ Virtual link — Connection is to a far-end router that is the endpoint of a virtual link.
LS Age	Time (in seconds) since LSA was originated
LS Seq	Sequence number of the LSA (used to detect older duplicate LSAs)
LSID	ID number of the router that originated the LSA
Metric	Cost of the link
Router ID	Originating router ID

**ip ospf linkStateData
network**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays network link state advertisements (LSAs) in the link state database.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`ip o l n`

3900
9300

Important Consideration

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
Area ID	Area ID (in the form n.n.n.n where 0 <= n <= 255) that corresponds to the OSPF area for which you want to view network LSA information	Valid area ID	0.0.0.0 (factory default), or current value
Area Mask	Subnet mask of OSPF area for which you want to view network LSA information	Valid area mask	0.0.0.0 (factory default), or current value
LSID	Link State ID: interface address of the designated router	Valid IP address	0.0.0.0 (factory default), or current value
LSID mask	Link State ID bit mask (Example: 255.0.0.0)	Link State ID bit mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data Network Display

Field	Description
Attached routers	List of routers that are fully adjacent to the designated router (DR); also the DR
LS Age	Time (in seconds) since the LSA was originated
LS Seq	Sequence number of the LSA (used to detect older duplicate LSAs)
LSID	Interface address (in the form of an IP address) of the designated router
Network mask	IP address mask for the network
Router ID	Originating router ID

**ip ospf linkStateData
summary****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays summary link state advertisements (LSAs) in the link state database.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o l s

Important Consideration

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
Area ID	Area ID (in the form n.n.n.n where 0 ≤ n ≤ 255) that corresponds to the OSPF area for which you want to view summary LSA information	Valid area ID	0.0.0.0 (factory default), or current value
Area mask	Subnet mask of the OSPF area for which you want to view summary LSA information	Valid area mask	0.0.0.0 (factory default), or current value
LSID	Link State ID: <ul style="list-style-type: none"> ■ For type 3 summary LSAs, this is the IP address of the destination network ■ For type 4 summary LSAs, this is the autonomous system boundary router's Router ID (in the form of an IP address) 	<ul style="list-style-type: none"> ■ For type 3 summary LSAs, a valid IP address ■ For type 4 summary LSAs, a valid router ID 	0.0.0.0 (factory default), or current value
LSID mask	Link State ID bit mask (Example: 255.0.0.0)	Link State ID bit mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data Summary Display

Field	Description
LS Age	Time (in seconds) since LSA was originated
LS Seq	Sequence number of the LSA (used to detect older duplicate LSAs)
LSID	<ul style="list-style-type: none">■ Type 3 — Destination network's IP address■ Type 4 — ASBR's OSPF router ID
Metric	Cost to reach the network
Network mask	<ul style="list-style-type: none">■ For Type 3 — destination network's IP address mask■ For Type 4 — Not used, must be 0 (--)
Router ID	Originating router ID

**ip ospf linkStateData
external****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays external network link state advertisements (LSAs) in the link state database.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**

ip o l e

3900
9300**Important Consideration**

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
LSID	Link State ID (in the form of the destination network's IP address)	Valid IP address	0.0.0.0 (factory default), or current value
LSID mask	Link State ID bit mask (Example 255.0.0.0)	Link State ID bit mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data External Display

Field	Description
Fwd Address	Forwarding address for data traffic to the advertised destination
LS Age	Time (in seconds) since LSA was originated
LS Seq	Sequence number of the LSA (used to detect older duplicate LSAs)
LSID	IP network number
Metric	Cost to reach advertised destination
Network Mask	IP address mask for the advertised destination
Router ID	Originating router ID
RouteTag	Not used by OSPF; these 32 bits may be used to communicate other information between boundary routers. Tag contents are defined by applications.
Type	<ul style="list-style-type: none"> ■ Type 1 — normal link state metric ■ Type 2 — metric is larger than any local link state path

**ip ospf neighbors
display**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays information about currently defined neighbors in an OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o n d

3900
9300

Fields in the IP OSPF Neighbors Display

Field	Description
Flags	Neighbor identification flags: <ul style="list-style-type: none"> ■ D — Dynamic neighbor ■ s — Static neighbor ■ BDR — Backup designated router ■ DR — Designated router Example: [S, BDR] + [D, DR] is a static neighboring backup designated router and a dynamic neighboring designated router
Indx	Interface index that corresponds to the interface to which a neighbor belongs
Neighbor Addr	Interface address of neighbor
Pri	Neighbor's OSPF router priority
ReqQ	Number of LSAs being requested from neighbor
Router ID	Neighbor's OSPF router ID
RxQ	Number of LSAs in local retransmit queue to the neighbor
State	Neighbor's adjacency: <ul style="list-style-type: none"> ■ Down — No recent data received from neighbor, connection is down. ■ Attempt — Only used on nonbroadcast networks. No recent data received from neighbor (will attempt to contact). ■ Init — Have recently seen Hello packet from neighbor; however, two-way communication has not been established. ■ Two-way — Bidirectional communication has been established. ■ ExStart — Taking initial step to create adjacency between neighboring routers. ■ Exchange — Database descriptions are being exchanged. ■ Loading — LSA databases are being exchanged. ■ Full — Neighboring routers are fully adjacent.
SumQ	Number of LSAs in LSA summary queue for the neighbor

ip ospf neighbors add *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- 9400

Adds a neighbor static IP address to an existing interface.

Valid Minimum Abbreviation

ip o n a

- 3900
- 9300

Important Consideration

- The system learns neighbor addresses dynamically on interfaces that support multicast routing. Define static neighbors only on nonmulticast interfaces.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the interface to which you want to add a neighbor	<ul style="list-style-type: none"> ■ Valid interface index number ■ ? (for a list of selectable indexes) 	First available (factory default), or current value
Static neighbor address	Address of neighbor that you want to define	Valid IP address on interface subnetwork	–

**ip ospf neighbors
remove**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes a static neighbor from an existing interface.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o n r

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the interface from which you want to remove a neighbor	<ul style="list-style-type: none"> ■ Valid interface index number ■ ? (for a list of selectable indexes) 	First available (factory default), or current value
Neighbor address	Address of neighbor that you want to remove	Valid IP address on interface subnetwork	–

ip ospf routerID *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- 9400

Sets the OSPF router ID.

Valid Minimum Abbreviation

ip o r

- 3900
- 9300

Important Considerations

- The OSPF router ID identifies the router to other routers within an autonomous system. Three types of router identifiers are available; all three take the form of an IP address:
 - **Default** — A unique ID that the system generates and uses as the default router ID
 - **Interface** — The index of an IP interface on the router
 - **Address** — An ID that you define in the form of an IP address
- OSPF routing must be inactive (disabled) before you can add or modify an OSPF router ID. To set the OSPF mode to `disabled`, see “ip ospf interface mode” earlier in this chapter. After you modify the router ID, you can set the OSPF mode to `enabled` on the interface
- The router ID must be unique from all other router IDs and ip interfaces in the autonomous system for OSPF to operate correctly. Choose the `default` setting to ensure unique router IDs.
- The resulting prompt depends on the router ID type that you choose.

Options

Prompt	Description	Possible Values	[Default]
Router ID type	Type of router identifier that you want to define	<ul style="list-style-type: none"> ■ default ■ interface ■ address 	default (factory default), or current value
IP interface	<i>For interface router ID type only.</i> Index number of IP interface to use as router ID.	<ul style="list-style-type: none"> ■ Valid IP interface ■ ? (for a list of selectable indexes) 	First available (factory default), or current value
Router ID	<i>For address router ID type only.</i> Identifier that is assigned to router in the form of an IP address 0.0.0.0 and 255.255.255.255 are invalid and will be rejected	User-defined router ID	Unique router ID generated by the system (factory default), or current value

IP OSPF Router ID Example (Interface Type)

```
Current OSPF router id = 0.43.66.0 (default)
Enter router ID type {default,interface,address|?} [default]: interface
Select IP interface {1-3|?}: 1
```

IP OSPF Router ID Example (Address Type)

```
Current OSPF router id = 24.23.11.23 (address)
Enter router ID type {default,interface,address|?} [address]: address
Enter router ID [24.23.11.23]: 101.89.2.4
```

ip ospf partition display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*
Displays OSPF memory allocation.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o pa d

Important Consideration

- See "ip ospf partition modify" later in this chapter for information on how OSPF memory allocation works and how to modify it.

3900
9300

Fields in the IP OSPF Partition Display

Field	Description
Current partition maximum size	OSPF memory partition upper limit as implemented at the last system reboot.
Configured partition maximum size	Last value that you entered, which will become the current partition maximum size after the next system reboot. <ul style="list-style-type: none"> ■ 0 means that OSPF has been set to use the system memory partition at the next reboot. ■ 1 means that OSPF has been set to use the default memory allocation scheme, deriving its partition size from the maximum size of the IP routing table at the next reboot. ■ Any other value that does not equal the current partition maximum size means that OSPF has been manually set to use a specific maximum partition size at the next reboot.
Allocated partition size	Module's current working memory. OSPF dynamically allocates memory in 100,000-byte chunks, up to the current partition maximum size.
OSPF is using the system partition	The administrator used the <code>ip ospf partition modify</code> command to set a partition value of 0. The OSPF protocol is using the system memory partition instead of its own partition, and there is no specified OSPF memory limit.

**ip ospf partition
modify****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Modifies the maximum memory that OSPF can allocate.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**`ip o pa m`**Important Considerations**

- There are three choices for memory allocation:
 - Have the system intelligently determine the maximum OSPF memory partition size (partition size = 1). This is the default.
 - Have OSPF be part of system memory, growing as needed and without limit (partition size = 0).
 - Configure the maximum OSPF memory partition size manually (partition size = 4096 - <maximum available memory>).
- You typically do not have to modify the OSPF memory allocation. However, if the `softRestarts` statistic shown by the `ip ospf statistics` option begins to climb, it means that OSPF is thrashing for memory and you must increase the maximum memory.



For a complete description of OSPF memory allocation, see the “OSPF Memory Partition” section in the OSPF chapter of the Implementation Guide.

- The partition size option that you enter takes effect after a system reboot.

Options

Prompt	Description	Possible Values	[Default]
New partition maximum size	Maximum memory size (in bytes) to allocate to OSPF system operations	<ul style="list-style-type: none"> ■ 4096 to <maximum available size> ■ 0 (to specify system memory partition) ■ 1 (to specify a size based on amount of memory and the maximum routing table size. On extended memory systems, this is 4,200,000.) 	1 (factory default), or current OSPF partition size

ip ospf
stubDefaultMetric
display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the stub default metric value for an area border router.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

ip o stu di

Important Considerations

- The stub default metric value determines if the router generates the default route into the stub areas of the network. This value applies to area border routers (ABRs) that have attached stub areas.



If a stub default metric is not defined, the router does not advertise a default route into the attached stub area.

- By default, the `stub default metric` is not defined.

Field in the IP OSPF Stub Default Metric Display

Field	Description
Stub default metric	Currently defined OSPF stub default metric

**ip ospf
stubDefaultMetric
define**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines the stub default metric value for an OSPF area border router.

Valid Minimum Abbreviation

`ip o stu de`

Important Considerations

- The stub default metric value determines if the router generates the default route into the stub areas of the network. This value applies to area border routers (ABRs) that have attached stub areas.



If a stub default metric is not defined, the router does not advertise a default route into the attached stub area.

- By default, the `stub default metric` is not defined.

Options

Prompt	Description	Possible Values	[Default]
Stub default metric	Stub default metric value to define for the area border router. Higher numbers are slower.	1 – 65535	Current stub default metric

✓ 3500
✓ 9000
9400

3900
9300

ip ospf
stubDefaultMetric
remove

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Disables the stub default metric on an OSPF area border router.

Valid Minimum Abbreviation

```
ip o stu r
```

Important Considerations

- The system removes the current stub default metric value immediately after you enter the command.
- The stub default metric value determines if the router generates the default route into the stub areas of the network. This value applies to area border routers (ABRs) that have attached stub areas.



If a stub default metric is not defined, the router does not advertise a default route into the attached stub area.

- By default, the `stub default metric` is not defined.

**ip ospf virtualLinks
summary****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays summary information about a virtual link.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**

ip o v su

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to display summary information	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–

Fields in the IP OSPF Virtual Links Summary Display

Field	Description
Dead Intvl	Number of seconds before the area border router's neighbors declare it down, when they stop hearing the router's Hellos
Hello Intvl	Length of time (in seconds) between Hello packets
Indx	Index number of the virtual link
Password	Password for the virtual link
Rxmit Intvl	Length of time (in seconds) between link state advertisement retransmissions
Target Router	End-point area border router where the virtual link terminates
Transit Area	Common area that the virtual link uses to reach the target router
Xmit Delay	Estimated number of seconds that it takes to transmit a link state update packet over the virtual link

**ip ospf virtualLinks
detail****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays detailed information about a virtual link.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**

ip o v det

Important Consideration

- This display also contains virtual link detail and neighbor information.

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to display detail information	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–

Fields in the IP OSPF Virtual Links Detail Display

Field	Description
Dead Intvl	Number of seconds before the area border router's neighbors declare it down, when they stop hearing the router's Hellos
Hello Intvl	Length of time (in seconds) between Hello packets
Indx	Index number of the virtual link
Password	Password for the virtual link
Rxmit Intvl	Length of time (in seconds) between link state advertisement retransmissions
Target Router	End-point area border router where the virtual link terminates
Transit Area	Common area that the virtual link uses to reach the target router
Xmit Delay	Estimated number of seconds that it takes to transmit a link state update packet over the virtual link

Fields in the IP OSPF Virtual Links Detail Display

Field	Description
Cost	Cost of sending a packet over the virtual link, expressed in the link state metric
Indx	Index number of the virtual link
Local Address	Address of the local router
Remote Address	Address of the remote router
State	State of the virtual link

Fields in the IP OSPF Virtual Links Neighbor Display

Field	Description
Indx	Index number for the interface to which a neighbor belongs
ReqQ	Number of LSAs that are being requested from the neighbor
RxQ	Number of LSAs that are in the local retransmit queue to the neighbor
State	Neighbor's adjacency: <ul style="list-style-type: none"> ■ Down — No recent data received from neighbor, connection is down. ■ Attempt — Only used on nonbroadcast networks. No recent data received from neighbor (will attempt to contact). ■ Init — Have recently seen Hello packet from neighbor; however, two-way communication has not been established. ■ Two-way — Bidirectional communication has been established. ■ ExStart — Taking initial step to create adjacency between neighboring routers. ■ Exchange — Database descriptions are being exchanged. ■ Loading — LSA databases are being exchanged. ■ Full — Neighboring routers are fully adjacent.
SumQ	Number of LSAs in LSA summary queue for the neighbor

**ip ospf virtualLinks
statistics****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays statistics that are associated with virtual links.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**

ip o v st

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to display statistics	<ul style="list-style-type: none"> ■ Valid interface index number ■ all ■ ? (for a list of selectable indexes) 	–

Fields in the IP OSPF Virtual Links Statistics Display

Field	Description
adjacencyDown	Number of times that OSPF adjacencies have gone down
adjacencyUp	Number of times that OSPF adjacencies have been formed
authError	Number of packets discarded due to OSPF authentication errors Interpretation: <ul style="list-style-type: none"> ■ A non-zero value is bad and means that packets from some OSPF routers are being discarded due to authentication errors. This statistic is incremented under the following circumstances: <ul style="list-style-type: none"> ■ If the OSPF packet authentication type is something other than simple password (that is, cryptographic authentication is not supported in the current implementation). ■ If the OSPF packet contains a password but the interface does not have a password configured. ■ If the OSPF packet has a simple password that does not match the password defined for the OSPF interface.
computeDR	Number of times that the designated router was computed
lsaXsumError	Number of LSA checksum errors that have been detected
mismatchAreaID	Number of interface area ID mismatches that have been detected
mismatchAreaType	Number of interface area type mismatches that have been detected

Field	Description
mismatchDead	<p>Number of router dead interval mismatches that were detected</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value is bad and means that some OSPF routers on the interface are configured with a different dead interval than this router. This prevents the router from becoming a neighbor with these other routers. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello packet is received and the dead interval it defines is different from the dead interval configured on the OSPF interface.
mismatchHello	Number of Hello packet interval mismatches that have been detected
mismatchMask	Number of subnet mask mismatches that have been detected
packetXsumError	Number of packet checksum errors since the interface has come up
receiveDD	<p>Number of database description packets that were received from valid OSPF neighbors.</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value is OK. <p>Database description packets are sent when forming adjacencies with valid neighbors. A large number of receiveDD packets in a network whose configuration has not changed could indicate that adjacencies are being torn down and reestablished.</p> <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF database descriptor packet from a valid OSPF neighbor is received.
receivedUnknown	Number of unknown LSAs that have been received

Field	Description
receiveError	<p>Number of general receive errors.</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value indicates that OSPF packets are being dropped and that this could be causing routing problems. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello packet is received and the packet length is too short. ■ When an OSPF Hello packet is received that has the same router ID as the router receiving the packet. ■ When an OSPF database descriptor packet is received and the packet length is too short. ■ When an OSPF link state request (LSR) packet is received and the packet length is too short. ■ When processing an LSR packet, if the area is not configured on the interface. ■ When an OSPF link state update (LSU) packet is received and the packet length is too short. ■ When processing an LSU packet, if there are more than 500 advertisements the packet is not processed. ■ When an OSPF link state acknowledgement (LSAck) packet is received and the packet length is too short. ■ When processing an LSAck packet, if the area described by the packet is not known by the router receiving the packet. ■ When processing any OSPF packet, if the packet length is less than the OSPF header length then it must have been truncated and the packet is dropped. ■ When an OSPF packet is received on an interface that is not running OSPF. ■ When an OSPF packet is received over a virtual link, but the virtual link is down or not configured. ■ When an OSPF packet is received (over a non-virtual link) from a source whose IP network does not match the IP network of the interface on which it was received. ■ When an OSPF packet is received on a Non-Broadcast Multiple Access network from an unknown neighbor. ■ When an OSPF packet is received whose version is not OSPF version 2.
receiveHello	Number of Hello packets that have been received
receiveLSAck	Number of LSA acknowledgments that have been received
receiveLSR	Number of LSA request packets that have been received

Field	Description
receiveLSU	Number of link state update packets that have been received
transmitDD	<p>Number of database description packets that were transmitted</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value is OK. <p>Database description packets are sent when forming adjacencies with valid neighbors. A large number in a network whose configuration has not changed could indicate that adjacencies are being torn down and re-established.</p> <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF database descriptor packet is transmitted.
transmitError	<p>Number of general transmit errors</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value indicates that an OSPF packet could not be sent either out a particular interface, or to a particular destination. This could prevent OSPF from running properly within the autonomous system and lead to routing problems. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello, LSU, or LSAck is being sent as a multicast packet on a non-broadcast multiple access network.
transmitHello	Number of Hello packets that have been transmitted
transmitLSAck	Number of LSA acknowledgments that have been transmitted
transmitLSR	Number of LSA request packets that have been transmitted
transmitLSU	Number of link state update packets that have been transmitted

ip ospf virtualLinks
define

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Creates a new virtual link to a destination router.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v def

Important Considerations

- All areas of an OSPF routing domain must connect to the backbone area. In cases where an area border router does not have direct, physical access to the backbone, you must configure a virtual link to act as a logical link to the backbone area.
- You can define up to 32 virtual links per router.

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Transit area	Area ID (in the form n.n.n.n where 0 <= n <= 255) through which the link is going	Currently defined area ID	–
Target router	ID of the target router, which is the router where the virtual link terminates	Valid IP address of OSPF area border router	–

**ip ospf virtualLinks
remove**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes a virtual link.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o v rem

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link that you want to remove	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–

ip ospf virtualLinks
areaID

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the transit area that is associated with a virtual link.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ip o v a

Options

- 3900
- 9300

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify a new area ID	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–
Target area	Area ID (in the form n.n.n.n where 0 <= n <= 255) of the transit area through which the virtual link must pass to reach the target router	ID of a currently defined area	Current value

**ip ospf virtualLinks
router**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the target router that is associated with a virtual link.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v ro

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify a new target router	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–
Target router	IP address of the new destination area border router where the virtual link terminates	Valid IP address of an OSPF area border router	Current value

ip ospf virtualLinks
delay

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the virtual link transmit delay, in seconds.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v del

Important Consideration

- The virtual link transmit delay must be consistent throughout the autonomous system.

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify the transmit delay	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–
Transmit delay	New virtual link transmit delay (in seconds)	1 – 65535 seconds	1 (factory default), or current value

**ip ospf virtualLinks
hello**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the virtual link Hello interval, in seconds.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

ip o v he

Important Considerations

- Hello packets inform other routers that the sending router is still active on the network.
- If a router does not send Hello packets for a period of time specified by the dead interval, the router is considered inactive by its neighbors.
- The virtual link Hello interval must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify the Hello interval	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–
Hello packet interval	Interval (in seconds) at which the area border router transmits Hello packets	1 – 65535 seconds	10 (factory default), or current value

ip ospf virtualLinks
retransmit

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the virtual link retransmit interval, in seconds.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v ret

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify the retransmit interval	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–
LSA retransmit time	Interval (in seconds) at which the area border router retransmits LSAs over the virtual link	1 – 65535 seconds	50 (factory default), or current value

**ip ospf virtualLinks
dead**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the virtual link dead interval, in seconds.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`ip o v dea`

Important Consideration

- Set the dead interval to the same value for all routers on the network.

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify the dead interval	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–
Dead interval	Maximum duration (in seconds) that neighbor routers wait for a Hello packet before they determine that the transmitting router is inactive	1 – 65535 seconds	40 (factory default), or current value

ip ospf virtualLinks
password

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets password security for a virtual link.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v p

Important Considerations

- Set the virtual link password to none to remove a previously assigned password.
- The password must be consistent throughout the autonomous system.

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify a password	<ul style="list-style-type: none"> ■ Valid IP interface index number ■ all ■ ? (for a list of selectable indexes) 	—
Virtual link password	Password for the specified virtual link	Up to eight ASCII characters	none (factory default), or current value

**ip ospf policy
summary**

Displays summary information about OSPF routing policies.

Valid Minimum Abbreviation

```
ip o p o s
```

Important Considerations

- Your system has one unified IP routing table. Routing policies allow you to control the flow of information among the network, the protocols, and the routing tables on your system.
- There are two classes of routing policies:
 - **Import policies** — Control which OSPF non-self-originated external routes are stored in the routing table. OSPF import policies control only what the local router uses. They do not affect the propagation of non-self-originated external routes to other routers.
 - **Export policies** — Used on OSPF boundary routers to control which self-originated external routing updates are placed in the link-state database for propagation over the network. In this way, export policies govern what other routers learn with regard to the local boundary router's self-originated information.
- The system tracks policies that you define in both OSPF and Routing Information Protocol (RIP), so the indexes that are assigned to your policies may have gaps. For example, if you have OSPF policies 1 and 2, and RIP policies 3 through 6, the next policy is 7.

Fields in the IP OSPF Policy Summary Display

Field	Description
Action	Action for the route (accept or reject)
Idx	Index number of the interface
Protocol	Protocol (for example, OSPF)
Route	Source network
Source	Source router
Type	Whether the policy is an import or export policy
Wt	Administrative weight (range of values: 1 through 16)

✓ 3500
✓ 9000
9400

3900
9300

ip ospf policy detail Displays summary and detailed information about OSPF routing policies.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

ip o po det

Important Considerations

- This display contains the summary information plus three additional fields: interface, metric, and ASEType.
- Your system has one unified IP routing table. Routing policies allow you to control the flow of information among the network, the protocols, and the routing tables on your system.
- There are two classes of routing policies:
 - **Import policies** — Control which OSPF non-self-originated external routes are stored in the routing table. OSPF import policies control only what the local router uses. They do not affect the propagation of non-self-originated external routes to other routers.
 - **Export policies** — Used on OSPF boundary routers to control which self-originated external routing updates are placed in the link-state database for propagation over the network. In this way, export policies govern what other routers learn with regard to the local boundary router's self-originated information.
- The system tracks policies that you define in both OSPF and Routing Information Protocol (RIP), so the indexes that are assigned to your policies may have gaps. For example, if you have OSPF policies 1 and 2, and RIP policies 3 through 6, the next policy is 7.

Fields in the IP OSPF Policy Detail Display

Field	Description
Action	Action for the route (accept or reject)
ASEType	Type of external metric — Type 1 or Type 2 — specified in the AS external link advertisement. OSPF boundary routers use Type 1 as default. Only applicable to export policies.
Index	Index number of the policy
Interface	Origin interface (only applicable when specifying direct as Origin Protocol)
Metric	Adjustment to the cost metric of routes that match the policy
Protocol	Origin protocol (for export policies only). Can also specify a direct or static route.

Field	Description
Route	Route against which the policy is applied
Source	Source router (only applicable to export policies that do not specify <code>direct</code> as Origin Protocol)
Type	Whether the policy is an <code>import</code> or <code>export</code> policy
Weight	Administrative weight (range of values: 1 through 16)

ip ospf policy define Defines import and export OSPF routing policies.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

```
ip o po def
```

Important Considerations

- The system assigns an index number to each policy and takes into account all route policies, Routing Information Protocol (RIP) and OSPF, that are set on the system.
- There are certain conditions associated with import and export policies. See the "OSPF Routing Policies" section in the OSPF chapter of your product's *Implementation Guide* for more information.
- Your system has one unified IP routing table. Routing policies allow you to control the flow of information among the network, the protocols, and the routing tables on your system.
- There are two classes of routing policies:
 - **Import policies** — Control which OSPF non-self-originated external routes are stored in the routing table. OSPF import policies control only what the local router uses. They do not affect the propagation of non-self-originated external routes to other routers.
 - **Export policies** — Used on OSPF boundary routers to control which self-originated external routing updates are placed in the link-state database for propagation over the network. In this way, export policies govern what other routers learn with regard to the local boundary router's self-originated information.
- You can set up an IP RIP or OSPF import or export policy to accept or advertise the default route, as long as the default route exists in the routing table. When you define a policy, you are always prompted for the route subnet mask after the route address, even though you specify the wildcard route address of 0.0.0.0.

Specify a route subnet mask as follows:

- If you want the wildcard subnet mask for all routes, use the default subnet mask (0.0.0.0).
- If you want the default route (not all routes), use 255.255.255.255.
- For more information about IP routing policies, see the *Implementation Guide* for your system.

Options

Prompt	Description	Possible Values	[Default]
Policy type	Type of policy	<ul style="list-style-type: none"> ■ import ■ export 	import
Origin protocols	For export policies only. Defines from which protocol the route originated	<ul style="list-style-type: none"> ■ direct ■ sta (static) ■ rip 	sta, rip
Source address	Source router from which the route was learned. Not applicable to the following: <ul style="list-style-type: none"> ■ Import policies ■ Export policies that define <code>direct</code> as the Origin Protocol 	Any valid IP address	0.0.0.0 (all)
Route address	Route IP address. Not applicable to export policies that define <code>direct</code> as the Origin Protocol.	Any valid IP address	0.0.0.0 (all)
Route subnet mask	Subnet mask for the route (for example, 255.255.0.0). Not applicable to export policies that define <code>direct</code> as the Origin Protocol.	Any valid subnet mask	0.0.0.0 (all)
IP interfaces	Index number of the interface for which you want to define a routing policy. Only applicable when specifying <code>direct</code> as the origin protocol when defining an export policy.	<ul style="list-style-type: none"> ■ Valid interface index ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Policy action	Accept or reject the route	<ul style="list-style-type: none"> ■ accept ■ reject 	accept
Metric adjustment	For <code>accept</code> conditions only, increases or decreases the converted route metric by the specified value. Options: <ul style="list-style-type: none"> + (add) - (subtract) * (multiply metric by value) / (divide metric by value) % (modulo, remainder of division operation as integer) 	0 – 65535 with or without options	0, which does not change the metric

Prompt	Description	Possible Values	[Default]
ASE type	Type of external metric that is used in the AS external advertisement (ASE), defined as: <ul style="list-style-type: none"> ■ Type 1 — External metric is directly comparable (without translation) to the link state metric. ■ Type 2 — External metric is larger than any link state path. 	<ul style="list-style-type: none"> ■ Type 1 ■ Type 2 	1
Administrative weight	Metric value for this policy. (Higher values have higher priority.)	1 – 16	1

OSPF Import Policy Conditions

Route (address/mask)	Action	Description
Specified route/mask	accept	Add specified non-self-originated external route with or without metric adjustments (+, -, *, /, %) to the routing table.
all (0.0.0.0)	accept	Add all non-self-originated external routes with or without metric adjustments (+, -, *, /, %) to the routing table.
Specified route/mask	reject	Do not add specified non-self-originated external route to the routing table.
all	reject	Do not add any external routes to the routing table; reject all non-self-originated external routes.

OSPF Export Policy Conditions

Protocol	Source Router	Route	Action	Description
RIP or static	Specified router or all routers	Specified route/mask	accept	Advertise in external LSAs specified RIP/static route from specified router with or without metric adjustments (+, -, *, /, %).
RIP or static	Specified router or all routers	all (0.0.0.0)	accept	Advertise in external LSAs all RIP/static routes from specified router with or without metric adjustments (+, -, *, /, %).
RIP or static	Specified router or all routers	Specified route/mask	reject	Do not advertise in external LSAs RIP/static routes from specified routers.
RIP or static	Specified router or all routers	all (0.0.0.0)	reject	Do not advertise in external LSAs any RIP/static route from specified routers.

Export Policy Conditions for Direct Routes

Protocol	Interface	Action	Description
Direct	Specified non-OSPF interface or All non-OSPF interfaces	accept	Advertise in external LSAs all direct routes off of specified interfaces.
Direct	Specified non-OSPF interface or All non-OSPF interfaces	reject	Do not specify in external LSAs any direct routes off of specified interfaces.

Example of Import Policy

```
Select menu option (ip/ospf/policy): define
Enter policy type (import,export) [import]: import
Enter route address [0.0.0.0]: 204.201.89.9
Enter route subnet mask [255.255.255.0]:
Enter policy action (accept,reject) [accept]: accept
Enter metric adjustment ([+,-,*,/,%]0-65535) [0]:
Enter administrative weight (1-16) [1]: 2
```

Example of Export Policy

```
Select menu option (ip/ospf/policy): define
Enter policy type (import,export) [import]: export
Enter origin protocols (dir,sta,rip|all|?) [dir,sta,rip]: sta
Enter source address [0.0.0.0]: 204.243.30.4
Enter route address [0.0.0.0]: 22.32.4.2
Enter route subnet mask [255.0.0.0]:
Enter policy action (accept,reject) [accept]: accept
Enter metric adjustment ([+,-,*,/,%]0-65535) [0]:
Enter ASE type (type1,type2) [type1]: 2
Enter administrative weight (1-16) [1]: 3
```

Example of Export Policy for a Directly Connected Interface

```
Select menu option (ip/ospf/policy): define
Enter policy type (import,export) [import]: export
Enter origin protocols (dir,sta,rip|all|?) [dir,sta,rip]: dir
Select IP interfaces (1|all|?) [1]:
Enter policy action (accept,reject) [accept]: accept
Enter metric adjustment ([+,-,*,/,%]0-65535) [0]: 3
Enter ASE type (type1,type2) [type1]: 2
Enter administrative weight (1-16) [1]: 4
```

ip ospf policy modify Modifies an existing OSPF routing policy.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o po m

Important Considerations

- The system assigns an index number to each policy and takes into account all route policies, Routing Information Protocol (RIP) and OSPF, that are set on the system.
- There are certain conditions associated with import and export policies. See the *Implementation Guide* for your system for more information.
- Your system has one unified IP routing table. Routing policies allow you to control the flow of information among the network, the protocols, and the routing tables on your system.
- There are two classes of routing policies:
 - **Import policies** — Control which OSPF non-self-originated external routes are stored in the routing table. OSPF import policies control only what the local router uses. They do not affect the propagation of non-self-originated external routes to other routers.
 - **Export policies** — Used on OSPF boundary routers to control which self-originated external routing updates are placed in the link-state database for propagation over the network. In this way, export policies govern what other routers learn with regard to the local boundary router's self-originated information.
- You can set up an IP RIP or OSPF import or export policy to accept or advertise the default route, as long as the default route exists in the routing table. When you define a policy, you are always prompted for the route subnet mask after the route address, even though you specify the wildcard route address of 0.0.0.0.

Specify a route subnet mask as follows:

- If you want the wildcard subnet mask for all routes, use the default subnet mask (0.0.0.0).
- If you want the default route (not all routes), enter 255.255.255.255.
- For more information about IP routing policies, see the *Implementation Guide* for your system.

Options

Prompt	Description	Possible Values	[Default]
Policy	Index number of the policy that you want to modify	<ul style="list-style-type: none"> ■ Valid policy index number ■ ? (for a list of selectable indexes) 	–
Origin protocols	For export policies only. Defines from which protocol the route originated	<ul style="list-style-type: none"> ■ direct ■ sta (static) ■ rip 	Current value
Source address	Source router from which the route was learned. Not applicable to the following: <ul style="list-style-type: none"> ■ Import policies ■ Export policies that define <code>direct</code> as the Origin Protocol 	Any valid IP address	Current value
Route address	Route IP address. Not applicable to export policies that define <code>direct</code> as the Origin Protocol.	Any valid IP address	Current value
Route subnet mask	Subnet mask for the route (for example, 255.255.0.0). Not applicable to export policies that define <code>direct</code> as the Origin Protocol.	Any valid mask	Current value
IP interfaces	Index number of the interface for which you want to define a routing policy. Only applicable when you specify <code>direct</code> as the origin protocol when defining an export policy.	<ul style="list-style-type: none"> ■ Valid IP interface index ■ all ■ ? (for a list of selectable indexes) 	Current value
Policy action	Accept or reject the route.	<ul style="list-style-type: none"> ■ accept ■ reject 	Current value

Prompt	Description	Possible Values	[Default]
Metric adjustment	For <code>accept</code> conditions only, increases or decreases the converted route metric by the specified value. Options: + (add) - (subtract) * (multiply metric by value) / (divide metric by value) % (modulo, remainder of division operation as integer)	0 – 16, with or without options	Current value
ASE type	Type of external metric used in the AS external advertisement (ASE), defined as: <ul style="list-style-type: none">■ Type 1 — External metric is directly comparable (without translation) to the link state metric.■ Type 2 — External metric is larger than any link state path.	<ul style="list-style-type: none">■ Type 1■ Type 2	Current value
Administrative weight	Metric value for this policy. (Higher values have higher priority.)	1 – 16	Current value

OSPF Import Policy Conditions

Route (address/mask)	Action	Description
Specified route/mask	accept	Add specified non-self-originated external route with or without metric adjustments (+, -, *, /, %) to the routing table.
All (0.0.0.0)	accept	Add all non-self-originated external routes with or without metric adjustments (+, -, *, /, %) to the routing table.
Specified route/mask	reject	Do not add specified non-self-originated external route to the routing table.
All	reject	Do not add any external routes to the routing table; reject all non-self-originated external routes.

OSPF Export Policy Conditions

Protocol	Source Router	Route	Action	Description
RIP or static	Specified router or all routers	Specified route/mask	accept	Advertise in external LSAs specified RIP/static route from specified router with or without metric adjustments (+, -, *, /, %).
RIP or static	Specified router or all routers	all (0.0.0.0)	accept	Advertise in external LSAs all RIP/static routes from specified router with or without metric adjustments (+, -, *, /, %).
RIP or static	Specified router or all routers	Specified route/mask	reject	Do not advertise in external LSAs RIP/static routes from specified router(s).
RIP or static	Specified router or all routers	all (0.0.0.0)	reject	Do not advertise in external LSAs any RIP/static route from specified router(s).

Export Policy Conditions for Direct Routes

Protocol	Interface	Action	Description
Direct	Specified non-OSPF interface or All non-OSPF interfaces	accept	Advertise in external LSAs all direct routes off of specified interfaces.
Direct	Specified non-OSPF interface or All non-OSPF interfaces	reject	Do not specify in external LSAs any direct routes off of specified interfaces.

IP OSPF Policy Modify Example

```
Select menu option (ip/ospf/policy): modify
Select policy {1|?} [1]:
Enter origin protocols (dir,sta,rip|all|?) [rip]:
Enter source address [0.0.0.0]:
Enter route address [0.0.0.0]:
Enter policy action (accept,reject) [accept]:
Enter metric adjustment ([+,-,*,/,%] 0-65535) [0]:
Enter administrative weight (1-16) [1]:
Enter ASE type (type1,type2) [type1]:
```

ip ospf policy remove Deletes OSPF routing policies.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o po r

Important Considerations

- The system assigns an index number to each policy that you define. This index number takes into account all route policies that are set on the system, Routing Information Protocol (RIP) and OSPF, so the assigned index may be higher than you expect.
- When you remove a policy, the associated index number is available for future use.

Options

Prompt	Description	Possible Values	[Default]
Policy index	Index number of the policy that you want to delete	<ul style="list-style-type: none"> ■ Valid policy index number ■ all ■ ? (for a list of selectable indexes) 	–

ip ospf statistics *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays general OSPF statistics.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`ip o sta`

3900
9300

Fields in the IP OSPF Statistics Display

Field	Description
extLsaChanges	Number of external LSA changes that have been made to the database
LSAsReceived	Number of link state advertisements that have been received
LSAsTransmitted	Number of link state advertisements that have been transmitted
memoryFailures	Number of nonfatal memory-allocation failures
recvErrors	Number of general receive errors
routeUpdateErrors	Number of nonfatal routing table update failures
softRestarts	Number of OSPF router soft restarts due to insufficient memory resources (implies a fatal memory-allocation failure). To fix this problem, use <code>ip ospf partition modify</code> to change the OSPF memory partition, add memory, or reconfigure the network topology to generate smaller OSPF databases.
SPFComputations	Number of shortest-path-first computations that have been made

This chapter provides guidelines and other key information about how to use the Internet Packet eXchange (IPX) protocol routing commands to route packets from your system to an external destination.

The IPX protocol is a NetWare LAN communications protocol that moves data between servers and workstation programs running on various network nodes. IPX is a User Datagram Protocol (UDP) that is used for connectionless communications. IPX packets are encapsulated and carried by Ethernet packet and Token Ring frames.

To route packets using the IPX protocol, you:

- 1 Define an IPX routing interface
- 2 Decide which IPX routing and server options you want to use
- 3 Enable IPX forwarding.

An IPX routing interface defines the relationship between an IPX virtual LAN (VLAN) and the subnetworks in the IPX network. Each routing IPX VLAN interface is associated with a VLAN that supports IPX. The system has one interface defined for each subnetwork that is directly connected to it. You must first define a VLAN, as described in Chapter 14, before you define an associated IPX VLAN interface.



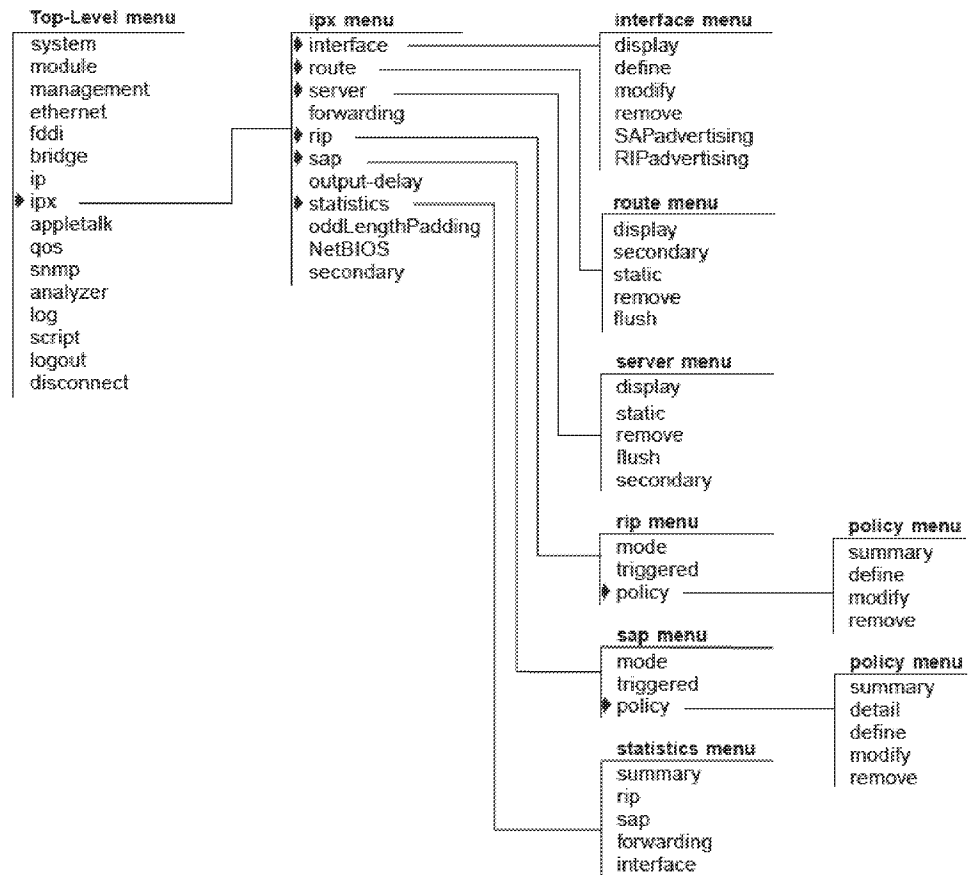
For more information about IPX, see the Implementation Guide for your system.



For the CoreBuilder® 9000, the commands in this chapter apply to Layer 3 switching modules only.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



ipx interface display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

3900
9300

Displays information about the IPX parameters and IPX interfaces that are configured on the system.

Valid Minimum Abbreviation

```
ipx i di
```

Important Considerations

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.

Fields in the IPX Interface Display

Field	Description
Format	Frame encapsulation format.
Index	System-assigned index number for the interface.
IPX address	Unique 4-byte network address.
State	Status of the IPX interface. It indicates whether the interface is available for communications (<code>up</code>) or unavailable (<code>down</code>).
Ticks	Number that the system uses to calculate route time. (A tick is an estimate of how long a packet takes to reach the network segment.) There are 18.21 ticks in a second. The possible values are 1 – 65534 and are defined as: <ul style="list-style-type: none"> ■ 1 = FDDI ■ 4 = Ethernet ■ 10+ = Serial Links
VLAN index	Index number of the VLAN that is associated with the IPX interface. When the system prompts you for this option, the menu identifies the available VLAN indexes.

ipx interface define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines an IPX interface.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

```
ipx i de
```

3900

9300

Important Considerations

- An IPX interface defines the relationships among an IPX virtual LAN (VLAN), the IPX router, and the IPX network. The IPX router has one IPX interface defined for each network than is directly connected to it.
- When you define an interface, you define the interface's IPX address, cost, format, and any associated IPX VLAN index.
- Before you define the IPX (routing) interface, you must specify a VLAN and select IPX, IPX-II, IPX-802.2, IPX 802.2 LLC, IPX-802.3, or IPX-802.2-SNAP as a protocol that the VLAN supports, as described in Chapter 14. (For routing, a VLAN can now support multiple protocols.)
- Unless your network has special requirements such as the need for redundant paths, assign a cost of 1 to each interface.
- The two Fiber Distributed Data Interface (FDDI) encapsulation formats correspond to the Ethernet 802.2 LLC and 802.3 SNAP encapsulation formats. If you select either of these Ethernet encapsulation formats, the corresponding FDDI encapsulation format is automatically selected for shared Ethernet and FDDI ports.

Options

Prompt	Description	Possible Values	[Default]
IPX network address	4-byte IPX address of the interface. The address must be unique within the network.	0x1 – 0xffffffffe	–
Ticks	Number that the system uses to calculate route time. (A tick is an estimate of how long a packet takes to reach the network segment.) There are 18.21 ticks in a second.	1 – 65534	1

Prompt	Description	Possible Values	[Default]
Frame format	Frame encapsulation format for the interface. IPX uses four Ethernet and two FDDI formats: Ethernet Type II, Novell 802.3 RAW, 802.2 LLC, and 802.3 SNAP. The FDDI formats are available with 802.2 and SNAP.	<ul style="list-style-type: none"> ■ Ethernet_II ■ 802.2 ■ 802.2 LLC ■ RAW_802.3 ■ SNAP ■ 802.3_SNAP 	–
VLAN Interface Index	Index number of the VLAN to associate with the IPX interface.	<ul style="list-style-type: none"> ■ A selectable VLAN interface ■ ? (to view a list of selectable indexes) 	–

IPX Interface Define Example

```

Select menu option: ipx interface define
Enter IPX Address (0x1-0xfffffffffe): 0x45468f30
Enter Ticks (1-65534) [1]:1
Enter Frame Format (Ethernet_II,802.2,Raw_802.3,SNAP): 802.2
Enter VLAN interface index {4|?} [4]: 4

```

ipx interface modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Changes the characteristics of an existing IPX interface.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

```
ipx i m
```

Important Considerations

- An IPX interface defines the relationships among an IPX virtual LAN (VLAN), the IPX router, and the IPX network. The IPX router has one IPX interface defined for each network that is directly connected to it.
- When you modify an interface, you can change the interface's IPX address, ticks, format, and the associated IPX VLAN index.
- Unless your network has special requirements (for example, a need for redundant paths), do not change the cost value of 1 that is assigned by default to each interface.

Options

Prompt	Description	Possible Values	[Default]
Index	Number associated with the interface that you want to modify.	<ul style="list-style-type: none"> ■ One or more selectable IPX interfaces ■ ? (to view a list of selectable interfaces) 	1 (if only 1 interface)
IPX network address	4-byte IPX address of the interface. The address must be unique within the network.	0x1 – 0xffffffe	Current address
Ticks	Number that the system uses to calculate route ticks. (A tick is an estimate of how long a packet takes to reach the network segment.) There are 18.21 ticks in a second.	1 – 65534 where: <ul style="list-style-type: none"> ■ 1 = FDDI ■ 4 = Ethernet ■ 10+ = Serial Link 	Current setting

Prompt	Description	Possible Values	[Default]
Frame format	Frame encapsulation format for the interface. IPX uses four Ethernet and two FDDI formats: Ethernet Type II, Novell 802.3 RAW, 802.2 LLC, and 802.3 SNAP. The FDDI formats are available with 802.2, SNAP, and 802.3/SNAP.	<ul style="list-style-type: none"> ■ Ethernet_II ■ 802.2 ■ 802.2 LLC ■ RAW_802.3 ■ SNAP ■ 802.3_SNAP 	Current format
VLAN interface index	Index number of the VLAN that is associated with the IPX interface.	<ul style="list-style-type: none"> ■ A selectable VLAN interface ■ ? (to view a list of selectable indexes) 	Current VLAN index

ipx interface remove***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

✓ 3500

✓ 9000

9400

3900

9300

Removes an IPX interface if you no longer perform routing on the ports that are associated with the interface.

Valid Minimum Abbreviation`ipx i r`**Options**

Prompt	Description	Possible Values	[Default]
Index	Index number for the interface that you want to remove	<ul style="list-style-type: none"> ■ One or more selectable IPX interface indexes ■ ? (to view a list of selectable indexes) 	1 (if only 1 interface)

**ipx interface
SAPadvertising**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Controls whether the system advertises IPX services.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ipx i s

Options

- 3900
- 9300

Prompt	Description	Possible Values	[Default]
IPX SAP advertising state	Whether the system advertises IPX services	<ul style="list-style-type: none"> ■ enable ■ disable 	disable

**ipx interface
RIPadvertising*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Controls whether the system advertises IPX routes.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ipx i r`**Options**

Prompt	Description	Possible Values	[Default]
IPX RIP advertising state	Whether the system advertises IPX services	<ul style="list-style-type: none"> ■ enable ■ disable 	disable

ipx route display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Displays the routing tables for the system. The routing tables include all configured routes.

Valid Minimum Abbreviation

```
ipx ro d
```

Important Considerations

- Your system maintains a table of routes to other IPX networks. You can:
 - Use the Routing Information Protocol (RIP) to exchange routing information automatically.
 - Make static entries in this table using the Administration Console.
- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.
- For a CoreBuilder 3500 system, the route table display shows the range for the routing table primary entries in the format $n - m$, where n is the current number of entries and m is the maximum number of primary entries.
- The maximum number of hops, or routers, that a packet can cross, is 16 (except NetBIOS packets, which can cross no more than 7 routers).

Options (3500 only)

Prompt	Description	Possible Values	[Default]
Start of address range	First address in a range for which you want to display routes	0x0 – 0xffffffff	0x0
End of address range	Last address in a range for which you want to display routes	0x0 – 0xffffffff	0xffffffff

Fields in the IPX Route Display

Field	Description
Address	Unique 4-byte network address of a segment in the system's routing table.
Age	Number of seconds that have elapsed since the last time the router sent a packet.
Hops	Number of hops, or the number of routers that must be crossed to reach the network segment.
Interface	System-assigned number for the interface.
Node	6-byte MAC address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
Ticks	Number of ticks, which is an estimate of time in seconds, that the packet takes to reach the network segment. There are 18.21 ticks in a second.

ipx route secondary Displays any secondary routes that are available.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

ipx ro se

Important Considerations

- To see entries for any secondary routes, you must:
 - Establish alternate paths to the same IPX network.
 - Enable the IPX secondary route/server option. See “ipx secondary” at the end of this chapter.
- A secondary route entry can replace a primary route entry when the primary route is removed from the routing table for any reason (for example, if the route reaches its age limit).
- For a CoreBuilder 3500 system, the route table display shows the range for the routing table primary entries in the format $n - m$, where n is the current number of entries and m is the maximum number of primary entries.

Fields in the IPX Secondary Route Display

Field	Description
Address	Unique 4-byte network address of a segment in the system's routing table.
Age	Number of seconds that have elapsed since the last time the router sent a packet.
Hops	Number of hops, or the number of routers that must be crossed to reach the network segment.
Interface	System-assigned number for the interface.
Node	6-byte MAC address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
Ticks	Number of ticks, which is an estimate of time in seconds, that the packet takes to reach the network segment. There are 18.21 ticks in a second.

ipx route static *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a static route.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ipx ro st`

Important Considerations

- Before you define static routes on the system, define at least one IPX interface. See “ipx interface define” earlier in this chapter for more details.
- Static routes remain in the routing table until you remove them or until you remove the corresponding interface.
- If an interface goes down, routes are temporarily removed from the routing table until the interface comes back up.
- Static routes take precedence over dynamically learned routes to the same destination. You can have a maximum of 32 static routes.

Options

Prompt	Description	Possible Values	[Default]
IPX network address	4-byte IPX address of the interface. The address must be unique within the network.	0x1 – 0xffffffffe	–
Hops	Number of hops, or number of routers that must be crossed to reach the network segment.	1 – 15	1
Interface number	Interface number to associate with the route. Depends on number of configured IPX interfaces.	<ul style="list-style-type: none"> ■ A selectable IPX interface number ■ ? (for a list of selectable IPX interfaces) 	–
Node address	6-byte MAC address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	A node address in the format xx-xx-xx-xx-xx-xx	–

IPX Static Route Example

```
Select menu option: ip route static
Enter IPX address (0x1-0xffffffffe): 0x44648f30
Enter Hops (1-15): 1
Enter interface number (1-32) [1]: 1
Enter node address: 08-00-3e-21-14-78
```


ipx route remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes a route from the IPX routing table.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

`ipx ro r`

3900

9300

Important Considerations

- The route is immediately deleted. You are not prompted to confirm the deletion.
- All servers that depend upon this route are removed from the server table, including static servers.

Options

Prompt	Description	Possible Values	[Default]
IPX network address	4-byte IPX network address	0x1 – 0xffffffff	–

ipx route flush *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes all dynamically learned routes from the IPX routing table.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

```
ipx ro f
```

3900

9300

Important Considerations

- All learned routes are immediately deleted. You are not prompted to confirm the deletion.
- All dynamic servers that depend on these routes are removed from the server table.

ipx server display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays the server table for the system to determine which servers are learned.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

```
ipx ser d
```

3900

9300

Important Considerations

- Your system maintains a table of servers that reside on other IPX networks. You can:
 - Use the Service Advertising Protocol (SAP) to exchange server information automatically.
 - Make static entries in this server table.
- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.
- For a CoreBuilder 3500 system, the route table display shows the range for the routing table primary entries in the format $n - m$, where n is the current number of entries and m is the maximum number of primary entries.

Options (3500 only)

Prompt	Description	Possible Values	[Default]
Service type	<p>Number for the type of service that the server performs.</p> <p>Enter up to 6 hex characters. For example, 0x4 = file server</p> <p>For more details, consult your Novell documentation.</p> <p>Use quotation marks (") around any string with embedded spaces.</p> <p>Use double quotes (" ") to enter an empty string.</p>	<ul style="list-style-type: none"> ■ * ■ 0x1 -0xffff 	*
Service name pattern	<p>Pattern for the service name.</p> <p>Use quotation marks (") around any string with embedded spaces.</p> <p>Use double quotes (" ") to enter an empty string.</p>	<ul style="list-style-type: none"> ■ * ■ Up to 79 alphanumeric characters 	*

Fields in the IPX Server Display

Field	Description
Age	Number of seconds that have elapsed since the last time a server in the table sent a packet.
Hops	Number of networks that must be crossed to reach the server. The maximum number is 15.
Interface	Index number of the interface.
Name	Name for the server that you define.
Network	4-byte IPX network address of the server.
Node	6-byte MAC address of the server that forwards packets to the segment.
Socket	2-byte socket address of the server that receives service requests.
Type	Type of service that the server provides. The IPX protocol defines various types of services. One common type is 0x4, which is for a file server. For more information on IPX type values, consult your Novell documentation.

ipx server static *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a static IPX server.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ipx ser st`

Important Considerations

- Static servers remain in the table until you remove them, until you remove the corresponding interface, or until you remove the route to the corresponding network address.
- A static server must have an IPX network address that corresponds to a configured interface or to a static route. If an interface goes down, any static servers on that interface are permanently removed from the server table until the interface comes back up.
- Static servers take precedence over dynamically learned servers to the same destination. You can have a maximum of 32 static servers.
- Before you define static servers on the system, first define at least one IPX interface. See "ipx interface define" earlier in this chapter for more details.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Interface index number for the server	<ul style="list-style-type: none"> ■ A selectable IPX interface index ■ ? (for a list of selectable IPX interfaces) 	–
Service type	Number for the type of service that the server performs	<ul style="list-style-type: none"> ■ * ■ 0x1 – 0xffff 	*
Service name	Service name of the server, up to 79 characters	<ul style="list-style-type: none"> ■ Any selectable service name ■ ? (for a list of selectable names) 	–
IPX network address	IPX network address of the server	0x0 – 0xffffffe	–
Socket value	Socket value of the server	0x0 – 0xffff	–
Node address	Node address of the server		–
Hops	Number of hops to the server	0 – 15	–

IPX Static Server Example

```
Enter Interface index {1|?} [1]: 1
Enter service type {0x1-0xFFFF}: 0x4
Enter service name {?}: gb201
Enter IPX address (0x0-0xffffffff): 0x8c14a228
Enter socket (0x0-0xffff): 0x8059
Enter node address : 00-00-2e-f3-56-02
Enter hops (0-15): 2
```

ipx server remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes a server from the IPX server table.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

`ipx ser r`

3900

9300

Important Consideration

- The server is immediately deleted. You are not prompted to confirm the deletion.

Options

Prompt	Description	Possible Values	[Default]
Service name	Service name of the server	<ul style="list-style-type: none"> ■ A selectable service name ■ ? (for a list of selectable names) 	
Service type	Number for the type of service that the server performs.	<ul style="list-style-type: none"> ■ * ■ 0x1 – 0xffff 	*

ipx server flush *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes all dynamically learned servers from the server table.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

`ipx ser f`

3900

9300

Important Consideration

- All learned servers are immediately deleted. You are not prompted to confirm the deletion.

ipx server secondary Displays any secondary servers that are available.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ipx ser se

Important Considerations

- To see entries for any secondary server, you must:
 - Establish alternate paths to the same IPX server.
 - Enable the IPX secondary route/server option. See “ipx secondary” at the end of the chapter.
- A secondary server entry can replace a primary server entry when the primary server is removed from the server table for any reason (for example, if the associated interface goes down, or the primary entry reaches its age limit).
- For a CoreBuilder 3500 system, the route table display shows the range for the routing table primary entries in the format $n - m$, where n is the current number of entries and m is the maximum number of primary entries.

Fields in the IPX Secondary Server Display

Field	Description
Age	Number of seconds that have elapsed since the last time a server in the table sent a packet.
Hops	Number of networks that must be crossed to reach the server. The maximum number is 15.
Interface	Index number of the interface.
Name	Name for the secondary server.
Network	4-byte IPX network address of the server.
Node	6-byte MAC address of the server that forwards packets to the segment.
Socket	2-byte socket address of the server that receives service requests.
Type	Type of service that the server provides. The IPX protocol defines various types of services. One type is 0x4, which is for a file server. For more information on IPX type values, consult your Novell documentation.

ipx forwarding *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Controls whether the system forwards or discards IPX packets.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ipx f

- 3900
- 9300

Important Considerations

- When you enable IPX forwarding, the system acts as a normal IPX router, forwarding IPX packets from one network to another when required.
- When you disable IPX forwarding, the system discards all IPX packets.

Options

Prompt	Description	Possible Values	[Default]
IPX forwarding state	Whether the system forwards or discards IPX packets	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled (factory default), or current value

ipx rip mode *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Selects the Routing Information Protocol (RIP) mode that is appropriate for your network.

Valid Minimum Abbreviation

`ipx ri m`

Important Considerations

- RIP allows the exchange of routing information on a NetWare network. IPX routers use RIP to create and maintain their dynamic routing tables.
- The system has three RIP modes:
 - **Off** — The system processes no incoming RIP packets and generates no RIP packets of its own.
 - **Passive** — The system processes all incoming RIP packets and responds to RIP requests, but it does not broadcast periodic or triggered RIP updates.
 - **Active** — The system processes all incoming RIP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered RIP updates.

Options

Prompt	Description	Possible Values	[Default]
RIP mode	Whether the system processes RIP packets	<ul style="list-style-type: none"> ■ off ■ passive ■ active 	disabled (factory default), or current value

ipx rip triggered *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Sets the RIP Triggered update mode, which dictates when the IPX protocol broadcasts newly learned routes.

Valid Minimum Abbreviation

`ipx ri t`

3900
9300

Important Considerations

- The system has two RIP triggered modes:
 - **Disabled** — Broadcasts IPX routes 3 seconds after learning them.
 - **Enabled** — Broadcasts IPX routes immediately after learning them.

Options

Prompt	Description	Possible Values	[Default]
Triggered update mode	Mode that determines when IPX broadcasts newly learned routes	<ul style="list-style-type: none"> ■ disabled ■ enabled 	enabled

**ipx rip policy
summary****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Display a list of IPX RIP (Routing Information Protocol) policies.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**`ipx ri p s`**Fields in an IPX RIP Policy Summary Display**

Field	Description
Idx	Index number of the IPX RIP policy.
Origin	Source of the route to which this policy applies. If the policy type is set to Export, the possible values of this parameter are RIP or Static. This parameter is not applicable if the policy type is set to Import.
Type	Import (apply the policy to received routes) or Export (apply the policy to advertised routes).
Route	One or more IPX network addresses where this policy applies.
Interface	One or more IP interfaces on this router associated with the RIP policy.
Source	6-byte MAC address of the router that can forward packets to the network. A source node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
Action	Whether this router accepts or rejects a route that matches the policy.
Metric	Value the system uses to increase or decrease a route metric. (This parameter is valid only if the Policy Action is set to Accept.)
Weight	Metric value of this policy.

ipx rip policy define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- 9400

Define a RIP (Routing Information Protocol) policy.

Valid Minimum Abbreviation

`ipx ri p d`

- 3900
- 9300

Important Considerations

- Every router maintains a table of current routing information in a routing table.
- Routing protocols receive or advertise routes from the network.
- Routing Policies control the flow of routing information between the network, the protocols, and the routing table manager.

Prompt	Description	Possible Values	[Default]
Type	Type of the policy: Import (apply the policy to received routes) or Export (apply the policy to advertised routes).	<ul style="list-style-type: none"> ■ Import ■ Export 	Import
Route Origin	Origin of the route to which this policy applies. This parameter is valid only if the policy Type is set to Export.	<ul style="list-style-type: none"> ■ Dir ■ Static ■ RIP ■ All 	All
Route address	Route to which this policy applies.	<ul style="list-style-type: none"> ■ 0x1-0ffffffe ■ All 	All
IP interfaces	One or more IP interfaces on this router associated with the RIP policy.	One or more IP interface numbers	All
Source node address	6-byte MAC address of the router that can forward packets to the network. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	<ul style="list-style-type: none"> ■ A node address in the format xx-xx-xx-xx-xx-xx ■ All 	All
Policy action	Whether this router accepts or rejects a route that matches the policy.	<ul style="list-style-type: none"> ■ Accept ■ Reject 	Accept

Prompt	Description	Possible Values	[Default]
Metric adjustment	Increase or decrease a route metric by a value that you specify. Specify an integer and an operand (+, -, *, /, %) to adjust the metric value. This parameter is valid only if the Policy Action is set to Accept.	<ul style="list-style-type: none"> ■ 0-16 ■ + (add) ■ - (subtract) ■ * (multiply) ■ / (divide) ■ % (modulo - remainder of integer division) 	0 (does not change the metric)
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same route. A higher value takes precedence over a lower value.	1 – 16	1

IPX RIP Policy Define Example

```

Select menu option (ipx/rip/policy): define
Enter policy type (import,export) [import]:export
Enter route origin (dir,static,rip,all) [all]:rip
Enter route address (0x1-0x1ffffffe[all]) [all]:all
Select IP interfaces (2[all?]) [all]:
Enter the source node address [all]:
Enter the policy action (accept, reject) [accept]: accept
Enter the metric adjustment ([+, -, *, /]0-16) [0]:
Enter the administrative weight (1-16) [1]:2

```

ipx rip policy modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

3900
9300

Modify an existing RIP (Routing Information Protocol) policy.

Valid Minimum Abbreviation

`ipx ri p m`

Important Considerations

- Every router maintains a table of current routing information in a routing table.
- Routing protocols receive or advertise routes from the network.
- Routing Policies control the flow of routing information between the network, the protocols, and the routing table manager.

Prompt	Description	Possible Values	[Default]
Policy	Index number of the policy you want to modify.	<ul style="list-style-type: none"> ■ 1 ■ ? (to view a list of selectable policies) 	1 (if only one policy)
Route Origin	Origin of the route to which this policy applies. This parameter is valid only if the policy Type is set to Export.	<ul style="list-style-type: none"> ■ Static ■ RIP ■ All 	All
Route address	IPX route to which this policy applies.	<ul style="list-style-type: none"> ■ 0x1-0xffffffe ■ All 	All
IP interfaces	One or more IP interfaces on this router associated with the RIP policy.	One or more IP interface numbers	All
Source node address	6-byte MAC address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	<ul style="list-style-type: none"> ■ A node address in the format xx-xx-xx-xx-xx-xx ■ All 	All
Policy action	Whether this router accepts or rejects a route that matches the policy.	<ul style="list-style-type: none"> ■ Accept ■ Reject 	Accept

Prompt	Description	Possible Values	[Default]
Metric adjustment	Increase or decrease a route metric by a value that you specify. Specify an integer and an operand (+, -, *, /, %) to adjust the metric value. This parameter is valid only if the Policy Action is set to Accept.	<ul style="list-style-type: none"> ■ 0-16 ■ + (add) ■ - (subtract)\ ■ * (multiply) ■ / (divide) ■ % (modulo - remainder of integer division) 	0 (does not change the metric)
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same route. A higher value takes precedence over a lower value.	1 – 16	1

IPX RIP Policy Modify Example

```

Select menu option (ipx/rip/policy): modify
Select policy {1|?}:1
Enter route origin (static,rip,all) [all]:rip
Enter route address (0x1-0x1ffffffe|all) [all]:
Select IP interfaces (2|all?) [all]:
Enter the source node address [all]:
Enter the policy action (accept, reject) [accept]:
Enter the metric adjustment ([+, -, *, /]0-16) [0]:
Enter the administrative weight (1-16) [1]:

```

ipx rip policy remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Remove an existing RIP (Routing Information Protocol) policy.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ipx ri p r

Options

Prompt	Description	Possible Values	Default
Policy	Index number of the policy you want to remove	<ul style="list-style-type: none"> ■ 1 ■ ? (to view a list of selectable policies) 	1 (if only one policy)

ipx sap mode *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Selects a Service Advertising Protocol (SAP) mode that is appropriate for your network.

Valid Minimum Abbreviation

`ipx sa m`

Important Considerations

- SAP provides routers and servers that contain SAP agents with a means of exchanging network service information.
- The system has three SAP modes:
 - **Off** — The system does not process any incoming SAP packets and does not generate any SAP packets of its own.
 - **Passive** — The system processes all incoming SAP packets and responds to SAP requests, but it does not broadcast periodic or triggered SAP updates.
 - **Active** — The system processes all incoming SAP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered SAP updates.

Options

Prompt	Description	Possible Values	[Default]
SAP mode	Whether the system processes SAP packets	<ul style="list-style-type: none"> ■ off ■ passive ■ active 	disabled (factory default), or current value

ipx sap triggered *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Sets the SAP Triggered Update mode, which dictates when the IPX protocol broadcasts newly learned SAP server addresses.

Valid Minimum Abbreviation

`ipx sa t`

3900
9300

Important Considerations

- The system has two SAP triggered modes:
 - **Disabled** — Broadcasts IPX SAP server addresses 3 seconds after learning them.
 - **Enabled** — Broadcasts IPX SAP server addresses immediately after learning them.

Options

Prompt	Description	Possible Values	[Default]
Triggered update mode	Setting for IPX SAP broadcast timing	<ul style="list-style-type: none"> ■ disabled ■ enabled 	enabled

**ipx sap policy
summary****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Display a list of IPX SAP (Service Advertising Protocol) policies.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**`ipx sa p s`3900
9300**Fields in an IPX SAP Policy Summary Display**

Field	Description
Idx	Index number of the IPX SAP policy.
Origin	Source of the service to which this policy applies. If the policy type is set to Export, the possible values of this parameter are SAP, Static, or All. This parameter is not applicable if the policy type is set to Import.
Type	Policy type. Import (apply the policy to received services) or Export (apply the policy to advertised services).
Name	Object name that assigned to the server.
Type	Service type, represented by a one-digit number. Refer to Novel documentation for a complete list of service types.
Network	IPX network address for the server, or All, which implies all routes.
Node	6-byte MAC address of the router that can forward packets to the network. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
Action	Whether this router accepts or rejects a route that matches the policy.

ipx sap policy detail *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
 ✓ 9000
 9400

Display information about IPX SAP (Service Advertising Protocol) policies.

Valid Minimum Abbreviation

`ipx sap p det`

3900
 9300

Fields in an IPX SAP Policy Detail Display

Field	Description
Idx	Index number of the IPX SAP policy.
Interface	Index number of the IP interface associated with this policy.
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same service. A higher value takes precedence over a lower value.

ipx sap policy define**For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

✓ 3500

✓ 9000

9400

3900

9300

Define a SAP (Service Advertising Protocol) policy.

Valid Minimum Abbreviation`ipx sa p def`**Important Considerations**

- Every router maintains a table of current configured services in a service table.
- The SAP running on the router receives and advertises services from the network.
- Service policies control the services in the service table and those that the router advertises.
- Novell defines several different service types using specific numbers for the server advertising the service. You enter a Novell service type when you define a SAP policy. Some of the most common service types are:

0x0004	File Server
0x0005	Job Server
0x0007	Print Server
0x0009	Archive Server
0x000A	Job Queue
0x0047	Advertising Print Server
0x0098	NetWare Access Server

For a complete list of Novell service types, consult your Novell documentation.

Options

Prompt	Description	Possible Values	[Default]
Policy Type	Type of the policy: Import (apply the policy to received services) or Export (apply the policy to advertised services).	<ul style="list-style-type: none"> ■ Import ■ Export 	Import
Service Origin	Origin of the service to which this policy applies. This parameter is valid only if the policy Type is set to Export.	<ul style="list-style-type: none"> ■ Static ■ SAP ■ All 	All

Prompt	Description	Possible Values	[Default]
Service Type	Number for the type of service that the server performs. Enter up to 6 hex characters. For example, 0x4 = file server For more details, consult your Novell documentation.	<ul style="list-style-type: none"> ■ 0x1 – 0xffff ■ All 	All
Server Name	Name of the server providing the services.	<ul style="list-style-type: none"> ■ Server name ■ All 	All
IPX Address	IPX network address of the network where the server resides.	<ul style="list-style-type: none"> ■ 0x0 – 0xffffffffe ■ All 	All
Node Address	6-byte MAC address of the router that can forward packets to the network. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	<ul style="list-style-type: none"> ■ A node address in the format xx-xx-xx-xx-xx-xx ■ All 	All
Interface Index	Index number of the IP interface associated with this policy.	<ul style="list-style-type: none"> ■ One or more interface numbers ■ All ■ ? (to view a list of selectable interfaces) 	All
Policy action	Whether this router accepts or rejects a service that matches the policy.	<ul style="list-style-type: none"> ■ Accept ■ Reject 	Accept
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same service. A higher value takes precedence over a lower value.	1 – 16	1

IPX SAP Policy Define Example

```
Select menu option (ipx/rip/policy): define
Enter policy type (import,export) [import]:
Enter service origin (static,sap,all) [all]:sap
Enter the service type (0x1-0x1ffff|all) [all]:0x0004
Enter the server name (?) [all]:
Enter the IPX address (0x0-0xffffffffe|all) [all]:
Enter the node address [all]:
Select interface index (2|all?) [all]:
Enter the policy action (accept, reject) [accept]: accept
Enter the administrative weight (1-16) [1]:2
```

ipx sap policy modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

3900
9300

Modify a SAP (Service Advertising Protocol) policy.

Valid Minimum Abbreviation

`ipx sa p m`

Important Considerations

- Every router maintains a table of current configured services in a service table.
- The SAP running on the router receives and advertises services from the network.
- Service policies control the services in the service table and those that the router advertises.
- Novell defines several different service types using specific numbers for the server advertising the service. You can change the Novell service type when you modify a SAP policy. Some of the most common service types are:

0x0004	File Server
0x0005	Job Server
0x0007	Print Server
0x0009	Archive Server
0x000A	Job Queue
0x0047	Advertising Print Server
0x0098	NetWare Access Server

For a complete list of Novell service types, consult your Novell documentation.

Options

Prompt	Description	Possible Values	[Default]
Policy	Index number of the policy you want to modify.	<ul style="list-style-type: none"> ■ 1 ■ ? (to view a list of selectable policies) 	1 (if only one policy)
Service Origin	Origin of the service to which this policy applies. This parameter is valid only if the policy Type is set to Export.	<ul style="list-style-type: none"> ■ Static ■ SAP ■ All 	All

Prompt	Description	Possible Values	[Default]
Service Type	Number for the type of service that the server performs. Enter up to 6 hex characters. For example, 0x4 = file server For more details, consult your Novell documentation.	<ul style="list-style-type: none"> ■ 0x1 – 0xffff ■ All 	All
Server Name	Name of the server providing the services.	<ul style="list-style-type: none"> ■ Server name ■ All 	All
IPX Address	IPX network address of the network where the server resides.	<ul style="list-style-type: none"> ■ 0x0 – 0xffffffffe ■ All 	All
Node Address	6-byte MAC address of the router that can forward packets to the network. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	<ul style="list-style-type: none"> ■ A node address in the format xx-xx-xx-xx-xx-x x ■ All 	All
Interface Index	Index number of the IP interface associated with this policy.	<ul style="list-style-type: none"> ■ One or more interface numbers ■ All ■ ? (to view a list of selectable interfaces) 	All
Policy action	Whether this router accepts or rejects a service that matches the policy.	<ul style="list-style-type: none"> ■ Accept ■ Reject 	Accept
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same service. A higher value takes precedence over a lower value.	1 – 16	1

IPX SAP Policy Modify Example

```
Select menu option (ipx/rip/policy): modify
Select policy {1|?}:1
Enter service origin (static,sap,all) [all]:sap
Enter the service type (0x1-0x1ffff|all) [all]:all
Enter the server name (?) [all]:
Enter the IPX address (0x0-0xffffffffe|all) [all]:
Enter the node address [all]:
Select interface index (2|all?) [all]:
Enter the policy action (accept, reject) [accept]: accept
Enter the administrative weight (1-16) [1]:2
```

ipx sap policy remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Remove an existing SAP (Service Advertising Protocol) policy.

Valid Minimum Abbreviation

`ipx sa p r`

Options

Prompt	Description	Possible Values	Default
Policy	Index number of the policy you want to remove	<ul style="list-style-type: none"> ■ 1 ■ ? (to view a list of selectable policies) 	1 (if only one policy)

ipx output-delay *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Sets the IPX output-delay option for RIP (Routing Information Protocol) and SAP (Service Advertising Protocol) packets. This option delays the updating of the RIP and SAP server information table.

Valid Minimum Abbreviation

ipx i o

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Output-delay mode	Whether you want to enable or disable the output-delay option	<ul style="list-style-type: none"> ■ enable ■ disable 	disable

**ipx statistics
summary****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays IPX summary statistics.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**`ipx st su`**Important Considerations**

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.

3900
9300**Fields in the IPX Statistics Summary Display**

Field	Description
Forwarded	Number of IPX packets that were forwarded
Fwd Received	Number of IPX packets that were received to be forwarded
Fwd Transmitted	Number of IPX forwarded packets that were successfully transmitted
Host Delivers	Number of IPX packets that were delivered to the IPX host's RIP and SAP applications
Host Dropped	Number of IPX packets to or from the IPX hosts's RIP and SAP applications that were dropped
Host Tx	Number of IPX packets from the IPX host's RIP and SAP applications that were successfully transmitted

ipx statistics rip *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays IPX RIP (Routing Information Protocol) statistics.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`ipx st r`

3900
9300

Important Considerations

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.

Fields in the IPX RIP Statistics Display

Field	Description
RIP Dropped	Number of IPX RIP packets that have been dropped
RIP Entries	Number of routes in the routing table (including local routes)
Routes Aged	Number of times the system marked a route entry unreachable, because it did not receive an update for that entry during the timeout period
RIP Received	Number of IPX RIP packets that have been received
RIP Requests	Number of IPX RIP requests that have been processed
RIP Responses	Number of IPX RIP responses that have been processed
RIP Transmitted	Number of IPX RIP packets that have been transmitted
Metric Changed	Number of times the metric changed on a route entry

ipx statistics sap**For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays IPX SAP (Service Advertising Protocol) statistics.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ipx st sa`**Important Considerations**

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.

Fields in the IPX SAP Statistics Display

Field	Description
SAP Dropped	Number of IPX SAP packets that have been dropped
SAP Entries	Number of servers in the server table
Servers Aged	Number of times the system marked a server entry unreachable, because it did not receive an update for that entry during the timeout period
SAP GNS Requests	Number of IPX SAP Get Nearest Service Requests that have been processed
SAP GNS Responses	Number of IPX SAP Get Nearest Service Responses that have been received
SAP Received	Number of IPX SAP packets that have been received
SAP Requests	Number of IPX SAP Requests that have been processed
SAP Responses	Number of IPX SAP Responses that have been processed
SAP Transmitted	Number of IPX SAP packets that have been transmitted
Metric Changed	Number of times the metric changed on a server entry

ipx statistics forwarding

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays IPX forwarding statistics.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ipx st f

Important Considerations

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered Updates are enabled.
 - Secondary route/server option is enabled.

- 3900
- 9300

Fields in the IPX Forwarding Statistics Display

Field	Description
Addr Errors	Number of IPX packets that were dropped that due to IPX address errors in network layer header
Forwarded	Number of IPX packets that were forwarded
Fwd Discards	Number of IPX packets to be forwarded that could not be forwarded
Fwd Received	Number of IPX packets that were received to be forwarded
Fwd Transmitted	Number of IPX forwarded packets that were successfully transmitted
Hdr Errors	Number of IPX packets that were dropped due to IPX Network layer header errors
Hop Count Errors	Number of IPX packets that were dropped due to exceeded maximum transport control
Host Delivers	Number of IPX packets that were delivered to the IPX host's RIP and SAP applications
Host In Discards	Number of IPX packets that were received for the IPX host's RIP and SAP applications that were dropped
Host Rx	Number of IPX packets that were delivered to the IPX host's RIP and SAP applications
Host Tx	Number of IPX packets that were transmitted from the IPX host's RIP and SAP applications

Field	Description
Host Tx Discards	Number of IPX packets from the IPX host's RIP and SAP applications that were dropped on transmission
Host Tx Request	Number of IPX packets from the IPX host's RIP and SAP applications to be transmitted
NetBIOS Max Hops	Number of IPX NetBIOS packets that exceeded the transport control maximum
NetBIOS Rx	Number of IPX NetBIOS packets that were received
NetBIOS Tx	Number of IPX NetBIOS packets that were transmitted
No Routes	Number of IPX packets that were dropped because the IPX route is unknown
Total Received	Number of IPX packets that were received
Tx Discards	Number of IPX packets that were forwarded but not successfully transmitted
Tx MTU Exceeded	Number of IPX packets that were forwarded but dropped because the MTU was exceeded

ipx statistics interface *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Displays IPX interface statistics.

Valid Minimum Abbreviation

`ipx st i`

3900
9300

Fields in the IPX Interface Statistics Display

Field	Description
Addr Errors	Number of IPX packets that were dropped due to IPX address errors in the network layer header
Forwarded	Number of IPX packets that were forwarded
Fwd Discards	Number of IPX packets to be forwarded that were not
Fwd Received	Number of IPX packets that were received to be forwarded
Fwd Transmitted	Number of IPX forwarded packets that were successfully transmitted
Hdr Errors	Number of IPX packets that were dropped due to IPX Network layer header errors
Hop Count Errors	Number of IPX packets that were dropped due to exceeded maximum transport control
Host In Discards	Number of IPX packets that were received for the IPX host's RIP and SAP applications that were dropped
Host Rx	Number of IPX packets that were received for the IPX host's RIP and SAP applications
Host Tx	Number of IPX packets that were transmitted from the IPX host's RIP and SAP applications
Host Tx Discards	Number of IPX packets from the IPX host's RIP and SAP applications that were dropped on transmission
Index	Index number that is assigned to the IPX interface
NetBIOS Max Hops	Number of IPX NetBIOS packets that exceeded the transport control maximum
NetBIOS Rx	Number of IPX NetBIOS packets that were received
NetBIOS Tx	Number of IPX NetBIOS packets that were transmitted
No Routes	Number of IPX packets that were dropped because the IPX route is unknown
Total Received	Number of IPX packets that were received
Tx Discards	Number of IPX packets that were forwarded but not successfully transmitted
Tx MTU Exceeded	Number of IPX packets that were forwarded but dropped because the MTU was exceeded

Field	Description
Routes Aged	Number of times the system marked a route entry unreachable, because it did not receive an update for that entry during the timeout period
Servers Aged	Number of times the system marked a server entry unreachable, because it did not receive an update for that entry during the timeout period
Rip Metric Changed	Number of times the metric changed on a route entry
Sap Metric Changed	Number of times the metric changed on a server entry

**ipx
oddLengthPadding**

- ✓ 3500
- ✓ 9000
- 9400

- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the compatibility mode for older network interface cards (NICs). This mode enables an interface to pad IPX packets that have an odd number of bytes. (Older NICs discard IPX packets that have an odd number of bytes.)

Valid Minimum Abbreviation

ipx od

Important Considerations

- This feature supports 10 MB switching modules only.
- If you use this feature, be careful to select only those interfaces that require odd-length padding. Enabling this feature for every interface slows network performance.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the interface for which you want to set the oddLengthPadding state	<ul style="list-style-type: none"> ■ A selectable IPX interface index ■ ? (for a list of selectable indexes) 	1 (if only 1)
IPX odd-length padding state	State for odd-length padding for the specified interface	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled

ipx NetBIOS Determines whether the system handles IPX Type 20 packet forwarding on a per-interface basis.

✓ 3500
9000
9400

Valid Minimum Abbreviation

`ipx n`

Options

Prompt	Description	Possible Values	[Default]
Interface index	index number of the interface for which you want to set the NetBIOS forwarding state	<ul style="list-style-type: none"> ■ One or more selectable IPX interface indexes ■ all ■ ? (for a list of selectable indexes) 	1 (if only 1)
IPX NetBIOS forwarding state	State for NetBIOS forwarding for the specified interface	<ul style="list-style-type: none"> ■ disabled ■ enabled 	enabled (factory default), or current value

IPX NetBIOS Example (3500)

```
Select menu option (ipx): netBIOS
Select interface index(es) (1-6|all|?): 1
Interface 1 - Enter state for NetBIOS packets
(disabled,enabled) [enabled]: disabled
```

ipx secondary Determines whether the system enables secondary routes and servers.

✓ 3500
9000
9400

3900
9300

Valid Minimum Abbreviation

ipx sec

Important Considerations

- This option allows the system to learn about secondary routes and secondary servers.
- With this option, a secondary route/server entry can replace a primary route/server entry when the primary route/server is removed from the routing/server table for any reason (for example, if the associated interface goes down, or if the primary entry reaches its age limit).
- For this option to have any effect, you must establish alternate paths to the same IPX network or server.
- After you enable the IPX secondary route/server option, you can display entries for any secondary routes or servers. (See "ipx route secondary" and "ipx server secondary" earlier in this chapter.)

Options

Prompt	Description	Possible Values	[Default]
IPX secondary route/server state	How to handle secondary routes and servers	<ul style="list-style-type: none"> ■ disabled ■ enabled 	enabled (factory default), or current value

IPX Secondary Example

```
Select menu option (ipx): secondary
Enter secondary route/server state (disabled,enabled)
[disabled]: enabled
```




APPLETALK

This chapter provides guidelines and other key information about commands that you can use to configure AppleTalk routing on your system. Configuring and managing AppleTalk routing involves these tasks:

- Administering AppleTalk interfaces
- Administering routes
- Administering the AARP cache
- Displaying the Zone Table
- Configuring forwarding
- Configuring checksum
- Enabling DDP Source Socket Verification
- Pinging an AppleTalk node
- Viewing AppleTalk statistics



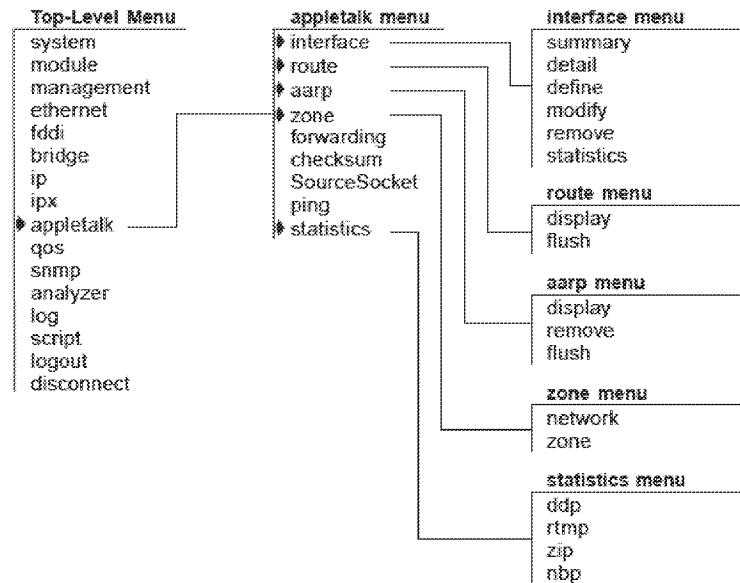
For more information about administering AppleTalk routing on your network, see the Implementation Guide for your system.



For the CoreBuilder® 9000, the commands in this chapter apply only to Layer 3 switching modules.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



**appletalk interface
summary**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays summary information for all AppleTalk interfaces.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ap i su

- 3900
- 9300

Fields in the AppleTalk Interface Summary Display

Field	Description
Address	AppleTalk interface address, which is based on the network range and the network node (Example: 203 01 . 7)
Index	Index number of the AppleTalk interface
Network range	Range of numbers that are assigned to the interface (Example: 203 01 – 203 10)
State	Status of the AppleTalk interface, which indicates whether the interface is available (enabled) or unavailable (down)
VLAN index	Index number of the virtual LAN (VLAN) that is associated with the AppleTalk interface

**appletalk interface
detail*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays detailed information for all AppleTalk interfaces.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**

ap i det

3900
9300**Fields in the AppleTalk Interface Detail Display**

Field	Description
Address	AppleTalk interface address, which is based on the network range and the network node. (Example: 20301.7)
Index	Index number of the AppleTalk interface
Network Range	Range of numbers that are assigned to the interface Example: (20301 – 20310)
Seed	Whether the interface is configured as a seed (y) or non-seed (n) interface
State	Status of the AppleTalk interface, that is, whether the interface is available (enabled) or unavailable (down)
VLAN index	Index number of the virtual LAN (VLAN) that is associated with the AppleTalk interface
Zone List	All zone names that are associated with the AppleTalk interface

appletalk interface define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*
 Defines an AppleTalk interface.

✓ 3500
 ✓ 9000
 9400

3900
 9300

Valid Minimum Abbreviation

ap i def

Important Considerations

- An AppleTalk interface defines the relationship between a virtual LAN (VLAN) and an AppleTalk network:
 - Every AppleTalk interface has one VLAN associated with it.
 - For routing purposes, you define a range of network numbers that are assigned to the AppleTalk interface. Example: 20301 – 20310
- You can configure the interface to be a seed or nonseed interface:
 - **Seed interface** — Initializes (“seeds”) the network with your configuration information. This information includes the network range and zone name list.
 - **Nonseed interface** — Listens for a seed router and then takes the zone and network range information from the first seed interface that it detects. After a nonseed interface obtains this information, it can participate in AppleTalk routing.
- Before you define the AppleTalk interface, you must define a VLAN and select AppleTalk as a protocol that the VLAN supports.
- Clients that have not been configured to use a particular zone use the default zone name.
- You can enter up to 16 zone names per interface.

Options

Prompt	Description	Possible Values	[Default]
Seed Interface	Whether an interface is configured as an AppleTalk seed (y) or non-seed interface (n).	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y (factory default), or current value
Start of network range	Start of the network range that is associated with the seed interface. <i>Seed interfaces only.</i>	1 – 65279	–
End of network range	End of the network range that is associated with the seed interface. <i>Seed interfaces only.</i>	1 – 65279	Value specified for start of network range, or current value

Prompt	Description	Possible Values	[Default]
Default zone name	User-defined default AppleTalk zone name. Clients that have not been configured to use a particular zone use the default zone name. <i>See interfaces only.</i>	Up to 32 ASCII characters	–
Zone name	AppleTalk zone that is associated with the interface. You are prompted to enter up to 15 additional zone names. <i>See interfaces only.</i>	<ul style="list-style-type: none"> ■ Up to 32 ASCII characters ■ q (to quit specifying zone names) 	–
VLAN interface index	Index number of the VLAN that you want to associate with the AppleTalk interface.	<ul style="list-style-type: none"> ■ Available valid VLAN index number ■ ? (for a list of available VLAN indexes) 	–

**appletalk interface
modify**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies an existing AppleTalk interface.

- ✓ 3500
- ✓ 9000
- 9400

- 3900
- 9300

Valid Minimum Abbreviation

ap i m

Important Considerations

- An AppleTalk interface defines the relationship between a virtual LAN (VLAN) and an AppleTalk network:
 - Every AppleTalk interface has one VLAN associated with it.
 - For routing purposes, you define a range of network numbers that are assigned to the AppleTalk interface. Example: 20301 – 20310
- You can configure the interface to be a seed or nonseed interface:
 - **Seed interface** — Initializes (“seeds”) the network with your configuration information. This information includes the network range and zone name list.
 - **Nonseed interface** — Listens for a seed router and then takes the zone and network range information from the first seed interface that it detects. After a nonseed interface obtains this information, it can participate in AppleTalk routing.
- Before you define the AppleTalk interface, you must define a VLAN and select AppleTalk as a protocol that the VLAN supports.
- Clients that have not been configured to use a particular zone use the default zone name.
- You can enter up to 16 zone names per interface.

Options

Prompt	Description	Possible Values	[Default]
Interface	Index number of the AppleTalk interface that you want to modify	<ul style="list-style-type: none"> ■ AppleTalk interface index number ■ ? (for a list of selectable indexes) 	–
Seed Interface	Whether you want to configure the interface as an AppleTalk seed (y) or nonseed interface (n)	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	Current value

Prompt	Description	Possible Values	[Default]
Start of network range	Start of the network range that is associated with the seed interface. <i>Seed interfaces only.</i>	1 – 65279	Current value
End of network range	End of the network range that is associated with the seed interface. <i>Seed interfaces only.</i>	1 – 65279	Current value
Default zone name	User-defined default AppleTalk zone name. Clients that have not been configured to use a particular zone use the default zone name. <i>Seed interfaces only.</i>	Up to 32 ASCII characters	Current value
Zone name	First AppleTalk zone that is associated with the interface. You are then prompted to enter up to 15 additional zone names. <i>Seed interfaces only.</i>	<ul style="list-style-type: none"> ■ Up to 32 ASCII characters ■ q (to quit specifying zone names and move on to the VLAN interface index prompt) 	Current value
VLAN interface index	Index number of the VLAN that you want to associate with the AppleTalk interface. When the system prompts you for a VLAN interface index, it indicates the available VLANs that you can associate with a new AppleTalk interface.	<ul style="list-style-type: none"> ■ Available valid VLAN index number ■ ? (for a list of selectable indexes) 	Current value
Interface down time	Number of minutes that you want to bring down the AppleTalk interface after you change zone information. This prompt appears only when you modify the zone information that is associated with the interface.	1 – 120 minutes	–

**appletalk interface
remove**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes an existing AppleTalk interface.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ap i r

Important Considerations

- You can specify a single interface, multiple AppleTalk interfaces, or all AppleTalk interfaces.
- If only one AppleTalk interface exists on the system, the interface is immediately removed after you enter this command.
- The system prompts you to select an interface number only if more than one AppleTalk interface exists on the system.

- 3900
- 9300

Options

Prompt	Description	Possible Values	[Default]
Interface	Index number of one or more interfaces that you want to remove	<ul style="list-style-type: none"> ■ One or more valid AppleTalk interface index numbers ■ ? (for a list of selectable indexes) ■ all 	–

appletalk interface statistics

✓ 3500
 ✓ 9000
 9400

3900
 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays statistics for each AppleTalk interface. You can specify a single AppleTalk interface, multiple interfaces, or all interfaces. If you have multiple interfaces and you do not specify one of them, the system prompts you to specify the appropriate interface index number.

Valid Minimum Abbreviation

ap i st

Important Consideration

- The display includes statistics for the AppleTalk Address Resolution Protocol (AARP), Datagram Delivery Protocol (DDP), Routing Table Maintenance Protocol (RTMP), Zone Information Protocol (ZIP), Name Binding Protocol (NBP), and AppleTalk Echo Protocol (AEP).

Fields in the AppleTalk Interface Statistics Display

Field	Description
aarpInProbes	Number of AARP probes that have been received
aarpInReqs	Number of AARP requests that have been received
aarpInResp	Number of AARP responses that have been received
aarpOutProbes	Number of AARP probes that have been sent
aarpOutReqs	Number of AARP requests that have been sent
aarpOutResp	Number of AARP responses that have been sent
ddpForwRequests	Total number of packets for which an attempt was made to forward them to their final destination
ddpInChecksumErrors	Number of DDP datagrams that were dropped because of a checksum error
ddpInLocals	Number of DDP datagrams for which this entity was the final DDP destination
ddpInReceives	Total number of packets that have been received, including those with errors
ddpInTooLongs	Number of input DDP datagrams that have been dropped because they exceeded the maximum DDP datagram size
ddpInTooShorts	Number of input DDP datagrams that have been dropped because the received data length was less than the data length that was specified in the DDP header, or the received data length was less than the length of the expected DDP header
ddpNoProtoHandlers	Number of DDP datagrams without protocol handlers
echoInReplies	Number of echo replies that have been received

Field	Description
echoInRequests	Number of echo requests that have been received
echoOutReplies	Number of echo replies that have been sent
echoOutRequests	Number of echo requests that have been sent
nbpInBroadcastReqs	Number of NBP broadcast requests that have been received
nbpInErrors	Number of NBP packets that have been received and rejected for any error
nbpInForwardReqs	Number of NBP forward requests that have been received
nbpInLookupReqs	Number of NBP lookup requests that have been received
rtmpInDataPkts	Number of RTMP data packets that have been received
rtmpInRequestPkts	Number of RTMP request packets that have been received
rtmpOutDataPkts	Number of good RTMP data packets that have been sent
rtmpRouteDeletes	Number of times that RTMP has deleted a route that was aged out of the table
zipAddressInvalids	Number of times that this entity had to broadcast a ZIP GetNetInfo reply because the GetNetInfo request had an invalid address
zipInErrors	Number of ZIP packets that have been received and rejected for any error
zipInExReplies	Number of ZIP extended replies that have been received
zipInGniRequests	Number of ZIP GetNetInfo request packets that have been received
zipInZipQueries	Number of ZIP queries that have been received
zipInZipReplies	Number of ZIP replies that have been received
zipOutGniReplies	Number of ZIP GetNetInfo reply packets that have been sent
zipOutInvalids	Number of ZIP GetNetInfo replies that have been sent with the indication that the previous client zone name was invalid

**appletalk route
display****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays AppleTalk routes that are listed in the system's routing table.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**

ap r d

Important Consideration

- Your system maintains a table of local and remote routes to all reachable AppleTalk networks. The Routing Table Maintenance Protocol (RTMP) automatically generates the routing table. RTMP defines rules for:
 - Information that is contained within each routing table
 - Exchanging information between routers so that the routers can maintain their routing tables

Fields in the AppleTalk Route Display

Field	Description
Distance	Distance in hops to the destination network
Interface	Interface that is used to reach the destination network
Network Range	Range of numbers that identify a network
Next Hop	Next hop Internet router to which the packet must be sent
State	Status of each route. One of the following: <ul style="list-style-type: none"> ■ good ■ suspect ■ bad ■ really bad

appletalk route flush *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes all dynamically learned AppleTalk routes from the routing table.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ap r f

3900

9300

Important Consideration

- The system deletes all dynamically learned AppleTalk routes immediately after you enter the command. You are not prompted to confirm the deletion.

**appletalk aarp
display**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the AppleTalk Address Resolution Protocol (AARP) cache.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ap a d

Fields in the AppleTalk AARP Display

Field	Description
AARP address	AppleTalk protocol address
Age (secs)	Age of the ARP entry (in seconds)
Interface	Index number of the interface on which the address was learned
MAC address	Hardware address that corresponds to the AppleTalk address

**appletalk aarp
remove**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes an AppleTalk Address Resolution Protocol (AARP) cache entry.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ap a r

Options

- 3900
- 9300

Prompt	Description	Possible Values	[Default]
AARP address	AARP address that you want to remove from the system's AARP cache	Any valid AARP address	–

appletalk aarp flush

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

Deletes all AppleTalk Address Resolution Protocol (AARP) entries from the system's AARP cache.

Valid Minimum Abbreviation

ap a f

3900

9300

Important Consideration

- The system deletes all AARP entries immediately after you enter the command. You are not prompted to confirm the deletion.

**appletalk zone
display network**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the AppleTalk Zone table, indexed by network numbers.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ap z d n

Important Considerations

- AppleTalk routers use the Zone Information Protocol (ZIP) to map network numbers to Zones.
- Each AppleTalk router maintains a Zone Information Table (ZIT), which lists the zone-to-network mapping information.

- 3900
- 9300

**appletalk zone
display zone**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the AppleTalk Zone table indexed by zones.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap z d z

Important Considerations

- AppleTalk routers use the Zone Information Protocol (ZIP) to map network numbers to Zones.
- Each AppleTalk router maintains a Zone Information Table (ZIT), which lists the zone-to-network mapping information.

appletalk forwarding

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables and disables AppleTalk Data Delivery Protocol (DDP) forwarding.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ap f

- 3900
- 9300

Options

Prompt	Description	Possible Values	[Default]
Forwarding state	Whether to enable or disable AppleTalk forwarding	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled (factory default), or current value

appletalk checksum *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Enables Data Delivery Protocol (DDP) checksum error detection for the AppleTalk protocol.

Valid Minimum Abbreviation

ap c

Important Considerations

- The AppleTalk protocol uses checksums to detect errors in data transmissions. A *checksum* totals all data bytes and adds the sum to the checksum field of the data packet. The receiving station computes a verification checksum from the incoming data and compares the new checksum with the value that is sent with the data. If the values do not match, the transmission contains an error.
- `Disabled` is the preferred setting. Enabling the checksum generation or verification significantly impacts the router's performance.

Options

Prompt	Description	Possible Values	[Default]
Checksum generation state	Whether to enable or disable generation of checksums for AppleTalk packets	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled (factory default), or current value
Checksum verification state	Whether to enable or disable verification of checksums for AppleTalk packets	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled (factory default), or current value

appletalk
sourceSocket

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables and disables AppleTalk Data Delivery Protocol (DDP) source socket verification.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ap so

Options

Prompt	Description	Possible Values	[Default]
source Socket	Whether to enable or disable source socket verification	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled (factory default), or current value

appletalk ping *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Pings an AppleTalk node using the AppleTalk Echo Protocol (AEP).

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ap p

Options

Prompt	Description	Possible Values	[Default]
Destination AARP address	AppleTalk node that you want to test for network connectivity	Valid AARP address	–

appletalk statistics
ddp

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays AppleTalk Datagram Delivery Protocol (DDP) statistics.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ap s d

- 3900
- 9300

Fields in the AppleTalk DDP Statistics Display

Field	Description
inBcastErrors	Number of dropped DDP datagrams for which the system was not their final destination and that were sent to the broadcast MAC address
inCsumErrors	Number of DDP datagrams that were dropped because of a checksum error
inDiscards	Number of DDP Datagrams that were discarded during routing
inForwards	Total number of packets that were forwarded, including those with errors
inLocals	Number of DDP datagrams for which an attempt was made to forward them to their final destination
inNoClients	Number of DDP datagrams that were dropped for unknown DDP types
inNoRoutes	Number of DDP datagrams that were dropped for unknown routes
inReceives	Total number of packets that were received, including those with errors
inShortDdps	Number of input DDP datagrams that were dropped because the system was not their final destination and their type was short DDP
inTooFars	Number of input datagrams that were dropped because the system was not their final destination and their hop count would exceed 15
inTooLongs	Number of input DDP datagrams that were dropped because they exceeded the maximum DDP datagram size
inTooShorts	Number of input DDP datagrams that were dropped because the received data length was less than the data length that was specified in the DDP header, or the received data length was less than the length of the expected DDP header
outLocals	Number of host-generated DDP datagrams

appletalk statistics
rtmp

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays AppleTalk Routing Table Maintenance Protocol (RTMP) statistics.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap s r

Fields in the AppleTalk RTMP Statistics Display

Field	Description
inDatas	Number of good RTMP data packets that were received
inOtherErrs	Number of RTMP packets that have been received and rejected for an error other than a version mismatch
inRequests	Number of good RTMP request packets that were received
inVersionErrs	Number of RTMP packets that have been received and rejected due to a version mismatch
outDatas	Number of RTMP data packets that were sent
outRequests	Number of RTMP request packets that were sent
routeDeletes	Number of times that RTMP deleted a route that was aged out of the table
routeEqChgs	Number of times that RTMP changed the Next Internet Router in a routing entry because the hop count that was advertised in a routing table was equal to the current hop count for a particular network
routeLessChgs	Number of times that RTMP changed the Next Internet Router in a routing entry because the hop count that was advertised in a routing table was less than the current hop count for a particular network
routeOverflows	Number of times that RTMP attempted to add a route to the RTMP table but failed because of lack of space

appletalk statistics
zip

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays AppleTalk Zone Information Protocol (ZIP) statistics.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap s z

3900
9300

Fields in the AppleTalk ZIP Statistics Display

Field	Description
inErrors	Number of ZIP packets that have been received and rejected for any error
inExReplies	Number of ZIP extended replies that have been received
inGniReplies	Number of ZIP GetNetInfo reply packets that have been received
inGniRequests	Number of ZIP GetNetInfo request packets that have been received
inLocalZones	Number of ZIP GetLocalZones requests packets that have been received
inObsoletes	Number of ZIP Takedown or ZIP Bringup packets that have been received
inQueries	Number of ZIP queries that have been received
inReplies	Number of ZIP replies that have been received
inZoneCons	Number of times that a conflict has been detected between this system's zone information and another entity's zone information
inZoneInvs	Number of times that this system has received a ZIP GetNetInfo reply with the zone invalid bit set because the corresponding GetNetInfo request had an invalid zone name
inZoneLists	Number of ZIP GetZoneLists requests packets that have been received
outAddrInvs	Number of times that this system had to broadcast a ZIP GetNetInfo reply because the GetNetInfo request had an invalid address
outExReplies	Number of ZIP extended replies that have been sent
outGniReplies	Number of ZIP GetNetInfo reply packets that have been sent out of this port
outGniRequests	Number of ZIP GetNetInfo packets that have been sent
outLocalZones	Number of transmitted ZIP GetLocalZones reply packets
outQueries	Number of ZIP queries that have been sent
outReplies	Number of ZIP replies that have been sent
outZoneInvs	Number of times that this system has sent a ZIP GetNetInfo reply with the zone invalid bit set in response to a GetNetInfo request with an invalid zone name
outZoneLists	Number of transmitted ZIP GetZoneList reply packets

appletalk statistics
nbp

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays AppleTalk Name Binding Protocol (NBP) statistics.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap s n

Fields in the AppleTalk NBP Statistics Display


Field	Description
inBcastReqs	Number of NBP Broadcast Requests that have been received
inErrors	Number of NBP packets that have been received and rejected for any error
inFwdReqs	Number of NBP Forward Requests that have been received
inLkupReplies	Number of NBP Lookup Replies that have been received
inLkupReqs	Number of NBP Lookup Requests that have been received

3900
9300



TRAFFIC POLICY

Chapter 22 Quality of Service (QoS) and RSVP



QUALITY OF SERVICE (QoS) AND RSVP

Quality of Service (QoS) and the *Resource Reservation Protocol (RSVP)* are advanced features that provide policy-based services. *Policy-based services* establish various grades of network services to accommodate the needs of different types of traffic (for example, multimedia, video, and file backups). QoS software relies on RSVP to provide admission control.

This chapter provides guidelines and other key information about how to configure QoS and RSVP in your system.

QoS and RSVP features include classifiers, controls, and RSVP parameters. Configure these features in the following order:

- 1 You first enter the command `qos`
- 2 to define how the system groups packets so that it can schedule them with the appropriate service level.
- 3 You then enter the command `qos control define` to assign rate limits and priorities to the packets that are associated with one or more of your classifiers. A classifier has no effect until you associate it with a control.

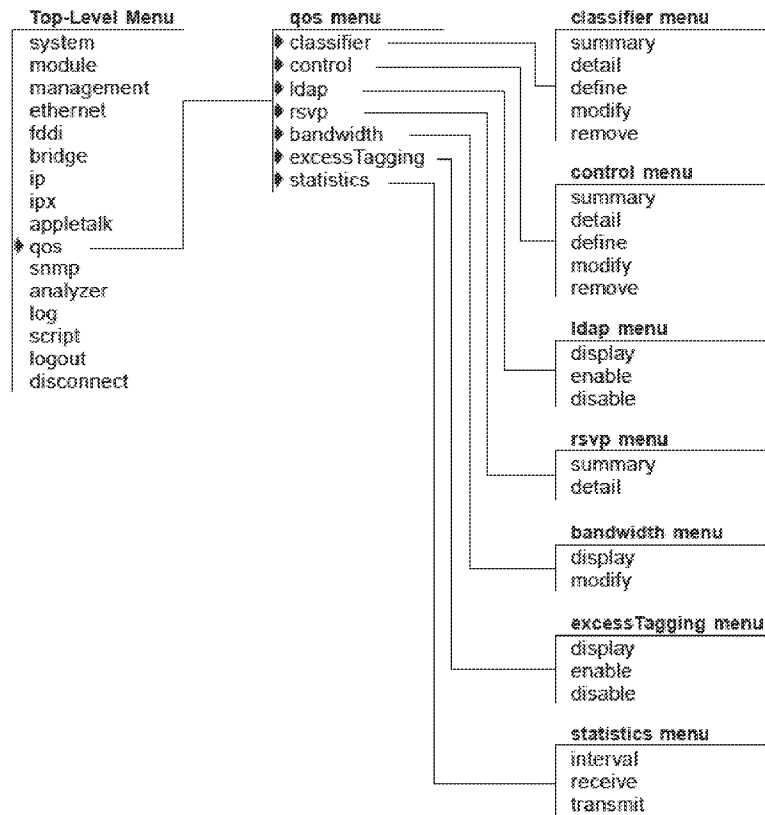
The system provides predefined classifiers and controls that are suitable for many configurations, or you can define your own classifiers, apply controls to the classifiers, and then decide whether to use RSVP. For more information about QoS and RSVP, see the *Implementation Guide* for your system.



For the CoreBuilder® 9000, the commands in this chapter apply only to Layer 3 switching modules.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



qos classifier summary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Displays summary information about the QoS classifiers on your system.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

q cl s

Fields in the QoS Classifier Summary Display

Field	Description
802.1p	For <i>nonflow</i> classifiers, IEEE 802.1p tag value
Cast	Cast type for the classifier: <ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: unicast, multicast, or all ■ <i>Nonflow</i> classifiers: unicast, multicast, broadcast, or all
Classifier	Number of the flow or nonflow classifier: <ul style="list-style-type: none"> ■ <i>Flow</i> classifiers in the range of 1 – 399 (Note: 20 and 23 are predefined.) ■ <i>Nonflow</i> classifiers in the range of 400 – 498 (Note: 401 – 407, 420, 430, 440, 450, 460, 470, 480, and 490 are predefined, but you can modify or remove them.)
Control	Control number that you assign to the classifier
Name	Name that you assign to the classifier
Protocol	Protocol type, if applicable, that is associated with the classifier/control: <ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: IP protocol type TCP, UDP, or all ■ <i>Nonflow</i> classifiers: TCP, IP, IPX, AppleTalk, or any

qos classifier detail *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays detailed information about one or more QoS classifiers.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q cl det

Options

Prompt	Description	Possible Values	[Default]
Classifier number	Number of the classifier for which you want detail information	<ul style="list-style-type: none"> ■ One or more numbers of configured classifiers ■ all ■ ? (for a list of selectable classifiers) 	–

Fields in the QoS Classifier Detail Display

Field	Description
802.1p	For nonflow classifiers, IEEE 802.1p tag value (any combination of priority tag values in the range 0 – 7)
Cast	The Cast type for the classifier: <ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: unicast, multicast, or all ■ <i>Nonflow</i> classifiers: unicast, multicast, broadcast, or all
Classifier	Number of the flow or nonflow classifier: <ul style="list-style-type: none"> ■ <i>Flow</i> classifiers in the range of 1 – 399 (Note: 20 and 23 are predefined.) ■ <i>Nonflow</i> classifiers in the range of 400 – 498 (Note: 401 – 407, 420, 430, 440, 450, 460, 470, 480, and 490 are predefined, but you can modify or remove them.)

Field	Description
Classifier – Filters (flow classifiers only)	Filters (address and port patterns): <ul style="list-style-type: none"> ■ Source IP address ■ Source IP address mask ■ Destination IP address ■ Destination IP address mask ■ Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port range
Destination Port range (flow classifiers only)	Beginning and end of the TCP or UDP destination port range
Source Port range (flow classifiers only)	Beginning and end of the TCP or UDP source port range
Classifier – Installed Flows (if flows exist)	Actual flows seen on the system, with the following data: <ul style="list-style-type: none"> ■ Port ■ Source IP address/source port ■ Destination IP address /destination port ■ Protocol type ■ Number of flow cache misses
Control	Control number that you assign to the control
Name	Name that you assign to the classifier
Protocol	Protocol type, if applicable, that is associated with the classifier and control

qos classifier define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a flow or nonflow classifier.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

q cl def

3900

9300

Important Considerations

- *Classifiers* define how the system groups packets so that it can schedule them with the appropriate service level. QoS supports flow and nonflow classifiers:
 - *Flow classifiers* apply to routed IP multicast and IP unicast packets. You can define up to 100 flow classifiers. Each filter (address and port pattern) in a flow classifier counts toward the limit.
 - *Nonflow classifiers* apply to bridged or routed traffic that is associated with a specific protocol (IP, TCP/IP, IPX, and AppleTalk) or to a custom protocol (EtherType or Destination Service Access Point/Source Service Access Point (DSAP/SSAP)). You can also use them to apply IEEE 802.1p tag values to forwarded frames. You can define up to 16 nonflow classifiers. All 16 nonflow classifiers are in use by default.
- The default classifier number is 499. You cannot remove or modify this default classifier. However, you can remove any of the predefined classifiers (for example, if you need another nonflow classifier). See “qos classifier remove” later in this chapter for more information.
- When you define a filter (address and port pattern) for a flow classifier, select a source and destination start and end port ranges that are as small as possible (for example, a single port). If the classifier applies to a wide range of Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports, you increase the amount of classified traffic on the system and consume valuable QoS resources.
- A classifier can have only one control applied to it.
- If you select `custom` when you define a nonflow classifier, you are prompted to select the protocol by EtherType or DSAP/SSAP. After you select a protocol, you are prompted to provide the hexadecimal ranges.



Depending on the number of VLANs defined, you can define a maximum of 3 custom protocols that can have controls applied to them. This limitation does not apply to non-controlled custom protocols.

Options

Prompt	Description	Possible Values	[Default]
Classifier number	Number of the flow or nonflow classifier in the range of 1 – 498	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: 1 – 399 (except 20 and 23, which are predefined flow classifiers) ■ <i>Nonflow</i> classifiers: 400 – 498, (except 401 – 407, 420, 430, 440, 450, 460, 470, 480, and 490. 401 – 407 are predefined nonflow classifiers with applied controls and IEEE 802.1p tag values of 1 – 7.) 	–
Classifier name	Name that you assign to the classifier	<ul style="list-style-type: none"> ■ Unique name with up to 32 characters (Use " around any string with embedded spaces. Use "" to enter an empty string.) ■ ? (for a list of selection criteria) 	–
Cast type	Cast type for the flow or nonflow classifier	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: unicast, multicast, or all ■ <i>Nonflow</i> classifiers: unicast, multicast, broadcast, or all ■ ? (for a list of selectable cast types) 	–
Protocol type	IP or other protocol type, if applicable, that you want to associate with the flow or nonflow classifier	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: IP protocol type with TCP, UDP, or all ■ <i>Nonflow</i> classifiers: TCP, IP, IPX, AppleTalk, custom, or any ■ ? (for a list of selectable protocol types) 	–
Source IP address	For <i>flow</i> classifiers only, IP address of the source	Up to 255.255.255.255	0.0.0.0 (factory default, wildcard match)

Prompt	Description	Possible Values	[Default]
Source IP address mask	For <i>flow</i> classifiers only, source IP address mask, or how many portions of the IP address you want to match (Example: 255.255.255.0 matches the first three portions of the specified IP address.)	Up to four portions (255.255.255.255)	0.0.0.0 (factory default)
Destination IP address	For <i>flow</i> classifiers only, destination IP address	Up to 255.255.255.255	0.0.0.0 (factory default, wildcard match)
Destination IP address mask	For <i>flow</i> classifiers only, destination IP address mask, or how many portions of the address you want to match	Up to four portions (255.255.255.255)	0.0.0.0 (factory default, wildcard match)
Start and end of TCP or UDP source port range	For <i>flow</i> classifiers only, start and end of the TCP or UDP source port range. The start value determines the end value.	<ul style="list-style-type: none"> ■ 0 – 65535 (start) ■ 2049 – 65535 (end) See "QoS Classifier Define Example (Flow Classifier)".	0 and 65535 (factory defaults)
Start and end of TCP or UDP destination port range	For <i>flow</i> classifiers only, start and end of the TCP or UDP destination port range. The start value determines the end value.	<ul style="list-style-type: none"> ■ 0 – 65535 (start) ■ 2049 – 65535 (end) See "QoS Classifier Define Example (Flow Classifier)".	0 and 65535 (factory defaults)
Additional filter (address/port pattern)	For <i>flow</i> classifiers only, additional source, destination, and port information for this classifier	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	no (factory default)
Custom protocol type (custom nonflow classifiers only)	For <i>nonflow</i> classifiers with the custom protocol type	<ul style="list-style-type: none"> ■ Ethertype ■ DSAP/SSAP 	–

Prompt	Description	Possible Values	[Default]
Custom protocol hexadecimal value (custom nonflow classifiers only)	Hex values for <i>nonflow</i> classifiers with the protocol custom type	<ul style="list-style-type: none"> ■ Ethertype hex value of 0x0 – 0xffe ■ DSAP hex value of 0x0 – 0xff Note: You cannot enter 0xaa - 0xaa ■ SSAP hex value of 0x0 – 0xff Note: You cannot enter 0xaa - 0xaa 	0x0 – 0x0
802.1p tag	For <i>nonflow</i> classifiers only, IEEE 802.1p tag values	<ul style="list-style-type: none"> ■ Any combination of priority tag values in the range of 0 – 7 ■ all ■ ? (for a list of possible values) 	–

Flow Classifier Procedure

To accept the default or current values that appear in brackets [], press Enter.

- 1 Enter a classifier number in the range of from 1 through 399.



Flow classifiers 20 and 23 are predefined for FTP and Telnet.

- 2 Enter the classifier name (a unique name of up to 32 characters).
- 3 Enter a cast type.

For a flow classifier, the options are *unicast*, *multicast*, and *all*.

- 4 Enter the IP protocol type of *TCP*, *UDP*, or *all*.
- 5 Enter the source IP address. The default value is 0.0.0.0.
- 6 Enter the source IP address mask. The default value is 0.0.0.0.
- 7 Enter the destination IP address.
- 8 Enter the destination IP address mask.
- 9 Enter the start of the TCP or UDP source port range, in the range of from 0 through 65535. The default is 0.
- 10 Enter the end of the TCP or UDP source port range using a value of up to 65535.

The value that you enter for the start of the range determines the default for the end of the range. The end value must be greater than or equal to the start value.



To avoid severely affecting applications using the network, select a port range that is as small as possible (for example, a single port).

- 11** Enter the start of the TCP or UDP destination port range, in the range of from 0 through 65535. The default is 0.
- 12** Enter the end of the TCP or UDP destination port range using a value of up to 65535. The end value must be greater than or equal to the start value.
- 13** At the prompt, specify whether you want any other filters (address and port patterns) with this classifier (yes or no). The default is no.

If you specify yes, the system prompts you for additional information, beginning with the source IP address.



Flow classifiers classify traffic only at the network layer and therefore affect only traffic that is being routed from one subnetwork to another.

QoS Classifier Define Example (Flow Classifier)

```
Select menu option (qos/classifier): define
Enter classifier number (1-498): 26
Enter classifier name {?}: IPFilter1
Select cast type (unicast,multicast|all|?): all
Select IP protocol type (TCP,UDP|all|?): all
Enter source IP address [0.0.0.0]:168.20.30.0
Enter source IP address mask [255.255.0.0]:255.255.255.0
Enter destination IP address [0.0.0.0]:192.1.0.0
Enter IP address mask [255.255.255.0]:255.255.0.0
Enter start of UDP source port range (0-65535) [0]:0
Enter end of UDP source port range (0-65535) [65535]:65535
Enter start of UDP destination port range (0-65535) [0]:0
Enter end of UDP destination port range (0-65535)
[65535]:65535
Enter another filter (yes,no) [no]: n
```


Nonflow Classifier Procedure

To accept the default or existing values that appear in brackets [], press Return.

- 1 Enter a classifier number in the range of from 400 through 498.
Numbers 401 through 407 are predefined nonflow classifiers with applied controls; numbers 420, 430, 440, 450, 460, 470, 480, and 490 are predefined nonflow classifiers without controls. If you have not removed any of the predefined nonflow classifiers, you need to remove them before you can define another nonflow classifier. (With the default classifier, there is a limit of 16 predefined nonflow classifiers.)
- 2 Enter the classifier name (a unique name of up to 32 characters long).
- 3 Enter a cast type.
For a *nonflow* classifier, the options are *unicast*, *multicast*, *broadcast*, and *all*.
- 4 Enter one or any protocols.
The options are *TCP/IP*, *IP*, *IPX*, *Appletalk*, *any*, or *custom*.
- 5 If you choose *custom*, enter the protocol type (*ethernet* or *DSAP/SSAP*).
 - For *ethernet* type enter the hexadecimal value.
 - For *DSAP/SSAP* type, enter the *DSAP* and *SSAP* hexadecimal values.
- 6 Enter one or all IEEE 802.1p tags. Specify any combination of values in the range of from 0 through 7, or *all*.

QoS Classifier Define Example (Nonflow Classifier)

```
Select menu option (qos/classifier): define
Enter classifier number (1-498): 481
Enter classifier name {?}: AppleBcast
Select cast type (unicast,multicast,broadcast|all|?):
broadcast
Select protocols {TCP/IP,IP,IPX,Appletalk,any,custom|?}:
Appletalk
Select IEEE 802.1p tag(s) (0-7|all|?): all
```

qos classifier modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- 9400

Modifies a previously defined classifier.

Valid Minimum Abbreviation

q cl m

- 3900
- 9300

Important Consideration

- If the classifier that you want to modify is associated with a control, you *must* remove the control before you can modify the classifier. See "qos classifier remove" later in this chapter for more information.

Options

Prompt	Description	Possible Values	[Default]
Classifier number	Number of the flow or nonflow classifier that you want to modify. Existing classifiers are shown in braces.	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: 1 – 399 (except 20 and 23, which are predefined flow classifiers) ■ <i>Nonflow</i> classifiers: 400 – 498, (except 401 – 407, 420, 430, 440, 450, 460, 470, 480, 490. 401 – 407 are predefined nonflow classifiers with applied controls.) ■ ? (for a list of selectable values) 	–
Classifier name	Name of the classifier that you want to modify	<ul style="list-style-type: none"> ■ Unique name with up to 32 characters (Use " around any string with embedded spaces. Use "" to enter an empty string.) ■ ? (for a list of selection criteria) 	Current name
Cast type	Cast type for the flow or nonflow classifier	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: unicast, multicast, or all ■ <i>Nonflow</i> classifiers: unicast, multicast, broadcast, or all 	Current cast type

Prompt	Description	Possible Values	[Default]
Protocol type	IP or other protocol type, if applicable, that is associated with the flow or nonflow classifier.	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: IP protocol type with TCP, UDP, or all ■ <i>Nonflow</i> classifiers: TCP, IP, IPX, AppleTalk, any, or custom ■ ? (for a list of selectable values) 	Current protocol type
Source IP address	For <i>flow</i> classifiers only, IP address of the source.	Up to 255.255.255.255	0.0.0.0 (factory default, wildcard match), or current value
Source IP address mask	For <i>flow</i> classifiers only, source IP address mask, or how many portions of the IP address you want to match. (Example: 255.255.255.0 matches the first three portions of the specified IP address.)	Up to four portions (255.255.255.255)	0.0.0.0 (factory default, wildcard match), or current value
Destination IP address	For <i>flow</i> classifiers only, destination IP address.	Up to 255.255.255.255	0.0.0.0 (factory default), or current value
Destination IP address mask	For <i>flow</i> classifiers only, destination IP address mask, or how many portions of the IP address you want to match.	Up to four portions (255.255.255.255)	0.0.0.0 (factory default), or current value
Start and end of TCP or UDP source port range	For <i>flow</i> classifiers only, start and end of the TCP or UDP source port range. Specify as small a range as possible. The start value determines the end value.	0 – 65535	0 and 65535 (factory defaults), or current values

Prompt	Description	Possible Values	[Default]
Start and end of TCP or UDP destination port range	For <i>flow</i> classifiers only, start and end of the TCP or UDP destination port range. Specify as small a range as possible. The start value determines the end value.	0 – 65535	0 and 65535 (factory defaults), or current values
Additional filters (address/port patterns)	For <i>flow</i> classifiers only, additional source, destination, and port information for this classifier. Each set of information counts toward the classifier limit.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	no (factory default)
Custom protocol type (custom nonflow classifiers only)	For <i>nonflow</i> classifiers with the custom protocol type.	<ul style="list-style-type: none"> ■ Ethertype ■ DSAP/SSAP 	–
Custom protocol hexadecimal value (custom nonflow classifiers only)	Hex values for <i>nonflow</i> classifiers with the custom protocol type.	<ul style="list-style-type: none"> ■ Ethertype hex value of 0x0 – 0xffff ■ DSAP hex value of 0x0 – 0xff ■ SSAP hex value of 0x0 – 0xff 	0x0 – 0x0
802.1p tag	For <i>nonflow</i> classifiers only, the IEEE 802.1p tag value	<ul style="list-style-type: none"> ■ Any combination of priority tag values in the range of 0 – 7 ■ all ■ ? (for a list of selectable values) 	Current value, if any

Procedure (Flow Classifier)

- 1 Enter the number of the classifier that you want to modify. The current numbers are shown in braces { }.
- 2 To modify the name, enter the new name for the classifier.

The name that is associated with the classifier number that you specified is shown in brackets.

- 3 To modify the cast type, enter a new cast type.
For a flow classifier, the options are `unicast`, `multicast`, and `all`.
To accept the default or current value that appears in brackets, press Enter.
- 4 To modify the IP protocol type, enter another IP protocol type (`TCP`, `UDP`, or `all`).
- 5 To modify the current source IP address, enter a new source IP address.
- 6 To modify the current source IP address mask, enter a new source IP address mask.
- 7 To modify the current destination IP address, enter a new destination IP address.
- 8 To modify the current destination IP address mask, enter a new destination IP address mask.
- 9 To modify the TCP or UDP source port range, enter the new start of the TCP or UDP port range (in the range of from 0 through 65535).
Limit the source port range as much as possible.
- 10 Enter the new end of the TCP or UDP source port range (in the range of from 0 through 65535).
- 11 To modify the TCP or UDP destination port range, enter the new start of the TCP or UDP port range (in the range of from 0 through 65535).
- 12 Enter the new end of the TCP or UDP destination port range (in the range of from 0 through 65535).
Limit the destination port range as much as possible.
- 13 At the prompt, specify whether you want any other address and port patterns (filters) with this classifier: `yes` or `no`; the default is `no`.
If you specify `yes`, the system prompts you for additional filtering information, beginning with the source IP address.



If you have several existing address and port patterns, you must specify all of them again during the modification process. Any address and port patterns that you do not reenter are deleted.

Nonflow Classifier Procedure

- 1 To modify the cast type, enter a new cast type.
For a nonflow classifier, the options are unicast, multicast, broadcast, and all
- 2 To modify the associated protocols, enter another protocol.
The options are TCP/IP, IP, IPX, Appletalk, any, or custom.
- 3 If you choose custom, select the protocol type (ethernet or DSAP/SSAP).
 - For the ethernet type, enter the hexadecimal value
 - For the DSAP/SSAP type, enter the DSAP and SSAP hexadecimal values
- 4 To modify the handling of IEEE 802.1p tags, enter the appropriate tags using a value in the range of 0 through 7, or enter all

QoS Classifier Modify Example (Flow Classifier)

```
Select menu option (qos/classifier): modify
Enter classifier number
{20,23,26,401-407,420,430,440, 450, 460, 470,480,490|?}:26
Enter classifier name {?} [IPfilter1]:
Select cast type (unicast,multicast|all|?)
[unicast,multicast]:
Select IP protocol type (TCP,UDP|all|?) [TCP,UDP]:
Enter source IP address [168.20.30.0]:
Enter source IP address mask [255.255.0.0]:
Enter destination IP address [192.1.1.0]:
Enter destination IP address mask [255.255.255.0]:
Enter start of TCP source port range (0-65535) [0]:
Enter end of TCP source port range (0-65535) [65535]:
Enter start of TCP destination port range (0-65535) [0]:
Enter end of TCP destination port range (0-65535) [65535]:
Enter another filter (yes,no) [no]: n
```

qos classifier remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Removes a previously defined classifier.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q cl r

Important Considerations

- If the classifier that you want to remove is associated with a control, you *must* remove the control before you can remove the classifier. See “qos control remove” later in this chapter for more information.
- When you enter the command, specify the number that represents the classifier that you want to remove, or specify ? to view the selectable classifiers.

Options

Prompt	Description	Possible Values	[Default]
Classifier number	Number for the classifier that you want to remove	<ul style="list-style-type: none"> ■ Any selectable classifier number ■ ? (for a list of selectable classifiers) 	–

QoS Classifier Remove Example (3500)

Select menu option: `qos classifier remove`

Enter classifier number

{20, 23, 26, 401-407, 420, 430, 440, 450, 460, 470, 480, 490|?}: 26

qos control summary *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays summary information about QoS controls.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

q co s

- 3900
- 9300

Fields in the QoS Control Summary Display

Field	Description
802.1p Tag	For controls for nonflow classifiers, the IEEE 802.1p tag value (0 – 7).
Classifiers controlled	Classifiers that this control affects.
Control number	Number of the control.
Control name	Name of the control.
Excess loss eligible	For receivePort or aggregate rate limit types, whether excess packets are loss eligible.
Excess service	For receivePort or aggregate rate limit types, the service level for excess packets.
Loss eligible	Whether conforming packets are loss eligible. If a packet is loss eligible, it can be dropped if the transmit queue for which it is destined exceeds its threshold.
Service	Service level for the conforming packets.

qos control detail *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays detailed information about the QoS controls that you specify.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q co det

Options

Prompt	Description	Possible Values	[Default]
Control number	Number of the control for which you want detail information	<ul style="list-style-type: none"> ■ One or more configured controls ■ all ■ ? (for a list of selectable controls) 	–

Fields in the QoS Control Detail Display

Field	Description
802.1p tag	IEEE 802.1p priority tag value (0 – 7) that is applied to forwarded frames. Can be defined for both flow and nonflow classifiers.
Burst	Burst size in KBytes.
Classifiers controlled	Classifiers that this control affects.
Control (number)	Number of the control.
Control name	Name that you assign to the control.
End time	Control end time.
Excess loss eligible	For receivePort or aggregate rate limit types, whether excess packets are loss eligible.
Excess service	For receivePort or aggregate rate limit type, service level for excess packets.
Limit	Rate limit in KBytes/sec or percentage.
Loss eligible	Whether conforming packets are loss eligible. A loss-eligible packet can be dropped if the transmit queue for which it is destined is over its threshold.
Ports	Receive ports for which you want to enable the rate limit.
Rate limits control	Number of the control that the rate limit affects.
Service	Service level for the conforming packets (high, best, low, or drop).

Field	Description
Source Port range	Beginning and end of the source port range.
Start time	Control start time
TCP drop control	Whether TCP drop control filtering is enabled.
Time control type	Time control type (<i>specific</i> , <i>daily</i> , <i>weekdays</i> , and so forth).
Type	Rate limit type, <i>none</i> (no rate limit), <i>receivePort</i> , or <i>aggregate</i> .

qos control define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a control for one or more existing classifiers.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q co def

Important Considerations

- A *control* can assign multiple rate limit values and an IEEE 802.1p priority tag value to the packets that are associated with one or more classifiers.
- The system prompts you according to the rate limit type that you select. You can only use one rate limit type (*none*, *receivePort*, or *aggregate*) per control. For a type of *receivePort* or *aggregate*, you can specify multiple rate-limit values for groups of ports or individual ports. The *aggregate* rate limit type can only be applied to flow classifiers.
- *Loss-eligible packets* are conforming packets that are discarded instead of queued when transmit queues back up beyond a threshold. You can specify whether conforming packets (as well as nonconforming excess packets) are loss eligible when you define the control. Marking packets loss eligible is useful for an intelligent discard of traffic in a congestion situation. Nonconforming excess packets are packets that exceed the specified rate limit.
- With the QoS timer control, you can configure QoS control sessions with starting and ending times (similar to using a VCR).

Options

Prompt	Description	Possible Values	[Default]
Control number	Number of the control. Control numbers 1 – 4 are predefined controls.	<ul style="list-style-type: none"> ■ 5 – 50 ■ ? (for a list of selectable values) 	1 (factory default)

Prompt	Description	Possible Values	[Default]
Control name	Name that you assign to the control. Predefined names are as follows: <ul style="list-style-type: none"> ■ Default/Best Effort (for control 1) ■ Background (for control 2) ■ Business Critical (for control 3) ■ Controlled Load (for control 4) 	<ul style="list-style-type: none"> ■ Unique name with up to 32 characters (Use " around any string with embedded spaces. Use "" to enter an empty string.) ■ ? (for a list of selection criteria) 	Default/Best Effort
Rate limit type	Type of rate limit: <ul style="list-style-type: none"> ■ none (no rate limit) ■ receivePort (a rate limit on the specified ports) ■ aggregate (the bandwidth for all ports chosen for the associated classifier). For flow classifiers only. 	<ul style="list-style-type: none"> ■ none ■ receivePort ■ aggregate 	none (factory default)
Service level	Service level for the conforming packets (a transmit priority that corresponds to a transmit queue). Drop causes the system to drop all traffic on all ports that are associated with the classifier and control.	<ul style="list-style-type: none"> ■ For rate limit receivePort or aggregate: high, best (best effort), or low ■ For a rate limit of none: high, best, low, or drop 	best (factory default)
Loss eligible	Whether conforming packets are loss-eligible. A loss-eligible packet can be dropped if the transmit queue for which it is destined exceeds its threshold.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	no (factory default)
Excess packet service	For receivePort or aggregate rate limit types, the service level for excess packets (packets that exceed the rate limit).	<ul style="list-style-type: none"> ■ high ■ best ■ low ■ drop 	best (factory default)
Excess loss eligible	For receivePort or aggregate rate limit types, whether excess packets are loss-eligible.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	yes (factory default)




Prompt	Description	Possible Values	[Default]
How rate limit is expressed	For receivePort or aggregate rate limit types, in KBytes/sec or percentage.	<ul style="list-style-type: none"> ■ KBytes/sec ■ percentage 	KBytes/sec (factory default)
Rate limit value	For receivePort or aggregate rate limit types, in KBytes/sec or percentage. 0 makes all packets excess packets.	<ul style="list-style-type: none"> ■ 0 – 65434 KBytes/sec ■ 0 – 100 percent 	–
Burst size	For receivePort or aggregate rate limit types, the maximum amount of data in Kbytes that you can transmit at the line rate before the transmission is policed.	16 – 8192 KBytes	Determined by your specified rate limit
Bridge ports	Receive ports for which you want to enable the rate limit. If you specify a subset of ports, you can specify multiple rate limit values. On the CoreBuilder® 9000, the list of ports includes the front-panel ports and any enabled backplane ports.	<ul style="list-style-type: none"> ■ Any subset of selectable ports ■ all ■ ? (for a list of selectable ports) 	Selectable ports
802.1p tag	IEEE 802.1p priority tag value to apply to forwarded frames (for both flow and nonflow classifiers).	<ul style="list-style-type: none"> ■ 0 – 7 ■ none ■ ? (for a list of selectable values) 	none (factory default)
Apply another rate limit?	If you specified a subset of available ports, whether you want to define another rate limit for other ports.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
TCP drop control enabled (flow classifiers only)	Whether one-way filtering is used so that drop packets establish a TCP connection.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
Start and end times	Whether you want to set starting and ending times for a control.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n

Prompt	Description	Possible Values	[Default]
Input time type	Type of time control that you want to establish. See Table 7 for a complete listing of input time type options.	<ul style="list-style-type: none"> ■ specific ■ daily ■ dayoftheweek ■ everydayoftheweek ■ weekdays ■ weekends ■ everyweekdays ■ everyweekends 	specific
Classifiers to be controlled	Classifiers for this control to affect. See "qos control summary" for a list of defined classifiers that are associated with controls.	Selectable classifiers (that is, those not already associated with a control)	—

Table 7 lists the options for the input time types. The key to the prompts are:

- mm/dd = month–day
- hh:mm = hour:minute

Table 7 Input Time Type Options

Input Time Type	Options
Specific (default)	Starting day (mm–dd) Starting time (hh:mm) Ending day (mm–dd) Ending time (hh:mm)
Daily	Starting day (mm–dd) Starting time (hh:mm) Ending day (mm–dd) Ending time (hh:mm)
	 <i>The Ending day and time cannot exceed 24 hours from the Starting day and time.</i>
Day of the week	Starting day (Monday = 1, Tuesday = 2, Wednesday = 3, Thursday = 4, Friday = 5, Saturday = 6, Sunday = 7) Starting time (hh:mm) Ending day (Monday = 1, Tuesday = 2, Wednesday = 3, Thursday = 4, Friday = 5, Saturday = 6, Sunday = 7) Ending time (hh:mm)
	 <i>The Ending day and time cannot exceed 24 hours from the Starting day and time.</i>
Every day of the week	Starting day (Monday = 1, Tuesday = 2, Wednesday = 3, Thursday = 4, Friday = 5, Saturday = 6, Sunday = 7) Starting time (hh:mm) Ending day (mm–dd) Ending time (hh:mm)
	 <i>The Ending day and time cannot exceed 24 hours from the Starting day and time.</i>
Weekdays	Starting time (hh:mm) Ending time (hh:mm)
Weekends	Starting time (hh:mm) Ending time (hh:mm)

Input Time Type	Options
Every weekday	Starting time (hh:mm) Ending time (hh:mm)
Every weekend	Starting time (hh:mm) Ending time (hh:mm)

Procedure

- 1 Enter a control number.

The valid range is 5 through 50, with the next available number as the default.

- 2 Enter a control name.

- 3 Enter the rate limit type: `none`, `receivePort`, or `aggregate`.

The default is `none`. To drop all conforming packets for a set of ports, use `receivePort` or `aggregate`, set the rate limit to 0, and specify the appropriate set of ports.



You can apply aggregate rate limits only to flow classifiers.

- 4 For the `receivePort` or `aggregate` limit type, enter the service level for conforming packets as `high`, `best`, or `low`.

For the `none` rate limit type, enter the service level for conforming packets as `high`, `best`, `low`, or `drop`.

The default is `best` (best effort).



If you use `drop`, the system drops all traffic on all ports for the classifier that is associated with the control. Ping packets are ICMP, not UDP/TCP, so they are not dropped.

- 5 Specify whether the conforming packets are loss eligible (`yes` or `no`).

The default is `no`.

- 6 If you have selected `receivePort` or `aggregate` for the rate limit type, you are prompted for the following information:

- a Enter the service level for excess packets (`high`, `best`, `low`, or `drop`). The default is `best`.

- b Specify whether excess packets are loss eligible (`yes` or `no`). The default is `yes`.

- c Specify how the rate limit is expressed (percentage of port bandwidth or `KBytes/sec`). `KBytes/sec` is the default.

- d** If you specified `KBytes/sec` for the rate limit, enter the value for the rate limit in `KBytes/sec` (0 through 65434).

If you specify that you want a percentage for the rate limit, specify the percentage in the range of from 0 through 100 percent. These numbers are rounded to the nearest 16 `KBytes/sec`. A value of 0 makes all packets excess packets.

- e** Enter the burst size in `KBytes` (16 through 8192, with the default value depending on your specified rate limit). The *burst size* is the maximum amount of data that you can transmit at the line rate before the transmission is policed.

- f** Specify the receive ports for which you want to enable the rate limit (specific bridge ports or all bridge ports).

If you apply the rate to only one or a subset of the bridge ports, you are prompted to specify whether you want to define another rate limit for another set of bridge ports. If you specify `yes`, you are prompted to enter another rate limit and burst size for another set of ports. This sequence of prompting continues until you specify `n`, meaning that you do not want to define another rate limit for another set of ports.



If the receive port is the anchor port for a trunk, the rate limit applies to each port that is associated with the trunk. For example, a rate limit of 1000 `KBytes` on a three-port trunk means that each port in the trunk has the 1000-`KByte` limit.

- 7** Enter an IEEE 802.1p tag value in the range of from 0 through 7 or `none` (the default) to apply to forwarded frames.
- 8** Specify whether drop packets used to establish a TCP connection (`yes`, `no`). The default is `no`.
- 9** Set the start and end time for the control (`yes`, `no`). The default is `no`.
 - a** If you specified a start and end time, enter the time type.

Time type selections are variations on days of the week and weekends or it can be specific day (or range of days) and time. See Table 7 for a complete listing of input time type options.
 - b** Enter the starting day and/or time.
 - c** Enter the ending day and/or time.

10 Enter the classifiers that are subject to this control.

The system displays the available classifiers in parentheses. If you select `aggregate` as the rate limit type, or if you said yes to the drop TCP connection packets option, only flow classifiers appear in parentheses.

QoS Control Define Example (3500)

This example shows a control for a nonflow classifier. Because the control has a rate limit of `none`, the system does not prompt you for information that applies to the other rate limit types.

```
Select menu option (qos/control): define
Enter control number {5-50|?} [5]:
Enter control name {?}: definetest
Enter rate limit type (none,receivePort,aggregate) [none]:
Enter service for conforming packets (high,best,low,drop)
[best]:
Are conforming packets loss eligible (yes,no) [no]:
Select IEEE 802.1p tag to apply to forwarded frames.
Tag {0-7|none|?} [none]:
Drop packets used to establish a TCP connection (yes,no)
[no]:
Set start and end time for the control (yes,no) [no]: yes
Enter input time type
(specific,daily,dayoftheweek,everydayoftheweek,weekdays,
weekends,everyweekdays,everyweekends) [specific]: weekdays
Enter the Qos control starting time (hh:mm): 09:00
Enter the Qos control ending time (hh:mm): 17:00
Select classifiers which are subject to this control.
Enter classifiers (20,23,420,430,440,450,4...: 450
```

qos control modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
 ✓ 9000
 9400

3900
 9300

Modifies the characteristics of a previously defined control (including controls 1 through 4, which the system provides by default).

Valid Minimum Abbreviation

q co m

Important Considerations

- The software prompts you according to the rate limit type that you select.
- If the existing control has a rate limit type of `receivePort` or `aggregate` with multiple rate limits, you can now change one rate limit without affecting the other defined rate limits.

Options

Prompt	Description	Possible Values	[Default]
Control number	Number of the control that you want to modify. Existing control numbers appear in braces. Control numbers 1-4 are predefined.	5 – 50	No default, but the next sequential number is automatically entered
Control name	Name of the control that you want to modify.	<ul style="list-style-type: none"> ■ Unique name with up to 32 characters (Use " around any string with embedded spaces. Use "" to enter an empty string.) ■ ? (for a list of selection criteria) 	Current name for specified control

Prompt	Description	Possible Values	[Default]
Rate limit type	Type of rate limit: <ul style="list-style-type: none"> ■ none (no rate limit) ■ receivePort (a rate limit on the specified ports) ■ aggregate (the bandwidth for all ports specified for the associated classifier) 	<ul style="list-style-type: none"> ■ none ■ receivePort ■ aggregate 	Current rate limit type
Service level	Service level for the conforming packets. Drop causes the system to drop all traffic on all ports that are associated with the classifier/control.	<ul style="list-style-type: none"> ■ For a rate limit of receivePort or aggregate: high, best (best effort), or low ■ For a rate limit of none: high, best, low, or drop 	Current service level
Loss eligible	Whether conforming packets are loss-eligible. A loss-eligible packet can be dropped if the transmit queue for which it is destined exceeds its threshold.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	Current value
Excess packets service	For receivePort or aggregate rate limit types, service level for excess packet.	<ul style="list-style-type: none"> ■ high ■ best ■ low ■ drop 	Current value
Excess loss eligible	For receivePort or aggregate rate limit types, whether excess packets are loss-eligible.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	Current value
How rate limit is expressed	For receivePort or aggregate rate limit types, format of the rate limit.	<ul style="list-style-type: none"> ■ KBytes/sec ■ percentage 	KBytes/sec
Rate limit value	For receivePort or aggregate rate limit types, number of Kbytes/sec or a percentage. 0 makes all packets excess packets.	<ul style="list-style-type: none"> ■ 0 – 65434 KBytes/sec ■ 0 – 100 percent 	–

Prompt	Description	Possible Values	[Default]
Burst size	For receivePort or aggregate rate limit types, maximum amount of data (in Kbytes) that you can transmit at the line rate before the transmission is policed.	16 – 8192 KBytes	Determined by your specified rate limit
Bridge ports	For receivePort or aggregate rate limit types, the receive ports for which you want to enable the rate limit.	<ul style="list-style-type: none"> ■ One or more selectable ports ■ all ■ ? (for a list of selectable ports) 	Current bridge ports
802.1p tag	IEEE 802.1p priority tag value that you want to apply to forwarded frames (for flow or nonflow classifiers).	<ul style="list-style-type: none"> ■ 0 – 7 ■ none ■ ? (for a list of selectable values) 	Current value
TCP drop control enabled (flow classifiers only)	Whether one-way filtering is used so that drop packets establish a TCP connection.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
Start and end times	Whether you want to set starting and ending times for a control.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
Input time type	Type of time control that you want to establish. See Table 7 for a complete listing of input time type options.	<ul style="list-style-type: none"> ■ specific ■ daily ■ dayoftheweek ■ everydayoftheweek ■ weekdays ■ weekends ■ everyweekdays ■ everyweekends 	specific
Classifiers controlled	Classifiers that this control affects. See “qos control summary” for a list of defined classifiers associated with controls.	Selectable classifiers (that is, those not already associated with a control)	Current classifier or classifiers for the control

Procedure

- 1 Enter the control number that you want to modify. The existing controls are displayed in braces { }.
- 2 To modify the name, enter the new name for the classifier.
The name that is associated with the specified control number appears in brackets [].
- 3 Enter the rate limit type (for example, `none`, `receivePort`, or `aggregate`).
The available values depend on how the control was defined; the current limit appears in brackets.
- 4 For the `receivePort` or `aggregate` rate limits, enter the service level for conforming packets as `high`, `best`, or `low`.
For the `none` rate limit, enter the service level for conforming packets as `high`, `best`, `low`, or `drop`. If you use `drop`, the system drops all traffic on all ports for the classifier that is associated with the control. The current value appears in brackets.
- 5 Specify whether the conforming packets are loss eligible (`yes` or `no`).
- 6 If you have selected `receivePort` or `aggregate` for the rate limit type, you are prompted for the following information:
 - a Enter the service level for excess packets (`high`, `best`, `low`, or `drop`).
 - b Specify whether excess packets are loss eligible (`yes` or `no`). Your current value is the default.
 - c Specify whether you want to modify the existing rate limits (`yes` or `no`).
If you enter `no`, the system maintains the existing values for all associated rate limits. If you enter `yes`, specify how the first rate limit should be expressed (percentage of port bandwidth or `KBytes/sec`). `KBytes/sec` is the default. If the control has multiple per-port rate limits, you can change one rate limit without affecting the others.
 - d If you specified `KBytes/sec` for the first (or only) rate limit, enter the value for the rate limit in `KBytes/sec` (0 through 65434).
If you specified percentage for the rate limit, specify the percentage in the range of from 0 through 100 percent.
 - e Enter the burst size in `KBytes` (in the range of from 16 through 8192). The default value depends on your specified rate limit.

- f** Specify the bridge ports for which you want to enable the new rate limit (for example, 1-13, or a11).

If you modify the rate limit and apply it to only one or a subset of the bridge ports, you are prompted to specify whether you want to modify or define another rate limit for another set of bridge ports. If you specify **yes**, you are prompted to enter another rate limit and burst size. This sequence of prompting continues until you specify **n**, meaning that you do not want to modify or define another rate limit for another set of ports. The rate limit applies only to those ports that you explicitly specified; any ports that you did not specify are not associated with your rate limit.
- 7** Select an IEEE 802.1p tag value in the range of from 0 through 7 or the value **none** to apply to forwarded frames.
- 8** Specify whether drop packets are used to establish a TCP connection (**yes**, **no**). The default is **no**.
- 9** Set the start and end time for the control (**yes**, **no**). The default is **no**.

 - a** If you specified a start and end time, enter the time type.

Time type selections are variations on days of the week and weekends or it can be specific day (or range of days) and time. See Table 7 for a complete listing of input time type options.
 - b** Enter the starting day and/or time.
 - c** Enter the ending day and/or time.
- 10** Enter the classifiers that are subject to this control. The system displays the associated classifiers in brackets. (If you select **aggregate** as the rate limit type, or select the **drop packets use to establish a TCP connection** option, the system displays only flow classifiers.)

QoS Control Modify Example (3500)

This example shows modifications to a predefined control (4) for a predefined classifier (405).

```
Select menu option: qos control modify
Enter control number {1-5}: 4
Enter control name {?} [Controlled Load]:
Interactive_Multimedia
Enter rate limit type (none, receivePort, aggregate) [none]:
receivePort
Enter service for conforming packets (high, best, low) [high]:
Are conforming packets loss eligible (yes, no) [no]:
Enter service for excess packets (high, best, low, drop) [low]:
drop
How should rate limit be expressed (percentage, KBytes/sec)
[KBytes/sec]:
Enter rate limit in KBytes/sec (0-65434): 2048
Enter burst size in KBytes (16-8192) [181]:
Select bridge ports (1-13|all|?) [1-13]:
Select IEEE 802.1p tag to apply to forwarded frames.
Enter IEEE 802.1p tag {0-7|none|?} [none]:
Drop packets used to establish a TCP connection (yes, no) [no]:
Do you want to modify/add the start and end time for the control (yes, no) [no]
Y
Do you want to have any time control (yes, no) [no]: y
Enter input time type (specific, daily, dayoftheweek, everydayoftheweek, weekdays,
weekends, everyweekdays, everyweekends) [specific]:
Enter the Qos Control starting day (mm-dd): 06-02
Enter the Qos control starting time (hh:mm): 09:00
Enter the Qos Control ending day (mm-dd): 06-02
Enter the Qos control ending time (hh:mm): 17:00
Select classifiers which are subject to this control.
Enter classifiers (20, 23, 404-407, 420, 430, 4... [404-407]: 405
```


qos control remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Removes a previously defined control.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

q co r

3900

9300

Important Consideration

- When you remove a control, the associated classifiers are no longer controlled and no longer have a set rate limit, service level, or 802.1p tag.

Options

Prompt	Description	Possible Values	[Default]
Control number	Number for the control that you want to remove	<ul style="list-style-type: none"> ■ One or more selectable control numbers ■ ? (for a list of the selectable controls) 	–

QoS Control Remove Example (9000 Layer 3)

```
CB9000@slot2.1 [12-E/FEN-TX-L3] (qos/control): remove
Enter control number {2-5|?}: 5
```

qos ldap display Displays Lightweight Directory Access Protocol (LDAP) status information.

✓ 3500
9000
9400

Valid Minimum Abbreviation

q 1 disp

Important Considerations

- When LDAP is enabled, displays server IP address and polling period.
- When LDAP is disabled, displays QoS, Resource Reservation Protocol (RSVP), and LDAP status.

3900
9300

Fields in the QoS LDAP Display

Field	Description
LDAP server address	The IP address of the LDAP server
Poll period	Selected poll period

qos ldap enable Enables QoS parameter directory services which are located on the Lightweight Directory Access Protocol (LDAP) server.

✓ 3500
9000
9400

Valid Minimum Abbreviation

q l e

3900
9300

Important Considerations

- An LDAP server must be configured.
- Before you enable LDAP, the LDAP server must have a directory group configured with QoS parameters in an *ldif* file.
- Parameter changes for a specific group may affect more than one system. If you know that a change will affect more than one system, disable LDAP to test the change. After you are sure you want the change, you can then enable LDAP.

Options

Prompt	Description	Possible Values	[Default]
Enable	Connects your system to the LDAP server	–	Disabled
Poll period		600 – 2000	–
LDAP server address	The IP address of the LDAP server you have configured	–	–
LDAP group name	Name of an LDAP entry on the LDAP server that indexes other entries containing QoS classifier and control information.	–	Wildcard

qos ldap disable Disables QoS parameter directory services, which are located on the Lightweight Directory Access Protocol (LDAP) server.

✓ 3500
9000
9400

Valid Minimum Abbreviation

q 1 disa

3900
9300

Important Considerations

- By default, LDAP is disabled.
- If LDAP is disabled, you do not receive automatic updates.

Options

Prompt	Description	Possible Values	[Default]
Disabled	Removes the connection to the LDAP server	–	Disabled

qos rsvp summary *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Displays summary Resource Reservation Protocol (RSVP) information when RSVP is enabled.

Valid Minimum Abbreviation

q r s

Fields in the QoS RSVP Summary Display

Field	Description
Excess loss eligible	Whether excess packets are loss-eligible.
Excess service	Service level for excess/policed traffic (<i>best</i> or <i>low</i>).
Per resv bandwidth	Largest reservation that RSVP attempts to install.
Policing option	When to drop excess packets. <i>Edge policing</i> causes excess packets to be dropped only at the edge (that is, when the traffic has not yet passed through any network device that has already performed policing for that flow). Options are <i>edge</i> , <i>always</i> , or <i>never</i> .
Total resv bandwidth	Admission control policy. RSVP begins to refuse reservations when the requested bandwidth on an output link exceeds the total reservable bandwidth.

qos rsvp detail *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
 ✓ 9000
 9400

Displays detailed RSVP information when RSVP is enabled.

Valid Minimum Abbreviation

q r de

Important Consideration

- If no flows are installed on the system or on a Layer 3 module, the command displays only the summary information.

3900
 9300

Options

Prompt	Description	Possible Values	[Default]
Level of RSVP information (when flows are installed)	If RSVP flows are available to report, the amount of RSVP information you want	<ul style="list-style-type: none"> ■ all ■ session ■ IP 	–

Fields in the QoS RSVP Detail Display

Field	Description
Excess loss eligible	Whether excess packets are loss-eligible.
Excess service	Service level for excess/policed traffic (<i>best</i> or <i>low</i>).
Per resv bandwidth	Largest reservation that RSVP attempts to install.
Policing option	When to drop excess packets. <i>Edge policing</i> causes excess packets to be dropped only at the edge (that is, when the traffic has not yet passed through any network device that has already performed policing for that flow).
Session	Session numbers, destination IP addresses and ports, protocols, number of senders, receivers, and RSVP reservations.
Session – receiver and session reservation	Port numbers, an RSVP style (<i>ST</i>) of fixed filter (<i>FF</i>), shared explicit (<i>SE</i>), or wildcard filter (<i>WF</i>), next hop addresses, LIH values, TTD values, bandwidth values, burst values, and filters.
Session – sender	Port numbers, source IP addresses, previous hop addresses, Logical Interface Handle (LIH) values, Time To Die (TTD) values, bandwidth values, burst size values, and output ports.
Session – installed flows	Actual flow that was installed on the system (shown in the last portion of the output).
Total resv bandwidth	Admission control policy. RSVP begins to refuse reservations when the requested bandwidth on an output link exceeds the total reservable bandwidth.

qos rsvp enable *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Enables RSVP on the system RSVP settings that you specify.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q r e

Important Considerations

- By default, RSVP is disabled.
- In general, when you enable RSVP, use the default settings.
- You are allowing RSVP to reserve this amount of bandwidth in the system. You can oversubscribe (over 100) and specify a value of up to 200.

Options

Prompt	Description	Possible Values	[Default]
Maximum total reservable bandwidth	Admission control policy. RSVP begins to refuse reservations when the requested bandwidth on an output link exceeds the total reservable bandwidth.	0 – 200 percent	50 (factory default)
Maximum per-reservation bandwidth	Largest reservation that RSVP attempts to install.	0 – 100 percent	50 (factory default)
Policing option	When to drop excess packets. <i>Edge policing</i> drops excess packets only at the edge (that is, when traffic has not yet passed through any network device that has already performed policing for that flow). <ul style="list-style-type: none"> ■ With <i>edge</i>, the system polices the flow when RSVP requests it. ■ With <i>always</i>, the system polices the flow regardless of whether RSVP requests it. ■ With <i>never</i>, the system never polices the flow even if RSVP requests it. 	<ul style="list-style-type: none"> ■ edge ■ always ■ never 	edge (factory default)

Prompt	Description	Possible Values	[Default]
Service level for excess /policed traffic	Service level for excess/policed traffic. Low is recommended. This setting applies to the excess traffic with the reserved bandwidth (that is, which queue it should be placed in).	<ul style="list-style-type: none">■ best■ low	low (factory default)
Excess Loss Eligible	Whether excess packets are loss-eligible	<ul style="list-style-type: none">■ yes■ no	no (factory default)

Procedure

- 1 Enter the maximum total reservable bandwidth, using a percentage of the output link (a value of from 0 through 200, with 50 as the default).
- 2 Enter the maximum per-reservation bandwidth, using a percentage of the output link (a value of from 0 through 100, with 50 as the default).
- 3 Enter the policing option (*edge*, *always*, or *never*, with *edge* as the default).
- 4 Enter the service level for excess/policed traffic (*best* or *low*, with *low* as the default).
- 5 Specify whether excess packets are loss eligible (*yes* or *no*, with *no* as the default).

qos rsvp disable *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Disables RSVP on the system.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

q r di

3900

9300

Important Considerations

- By default, RSVP is disabled.
- This command does not verify that RSVP has been disabled.

**qos bandwidth
display**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the link bandwidth as the ratio of bandwidth that is allocated to high priority traffic versus best effort traffic. Link bandwidth is the total link bandwidth less the bandwidth that RSVP and network control traffic use.

Valid Minimum Abbreviation

q b d

Important Consideration

- By default, 75 percent of bandwidth is allocated to high-priority traffic.

**qos bandwidth
modify**

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets how to weigh the high priority and best effort transmit queues, and sets RSVP bandwidth for the control queue. Low priority packets do not have bandwidth explicitly allocated.

Valid Minimum Abbreviation

q b m

Important Considerations

- When you enter the command, the system prompts you to enter the percentage of bandwidth to use for high-priority traffic on the output link.
- The value 75 specifies that three high-priority packets are transmitted for each best effort packet.
- The value 50 sets equal priority for high priority and best effort packets.
- The value 100 is strict prioritization; it allows best effort packets to be sent only when no high priority packets need to be sent.

Options

Prompt	Description	Possible Values	[Default]
Percentage of bandwidth	Percentage of bandwidth that you want to be used for high-priority traffic on the output link	0 – 100 percent	75

qos excessTagging display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays status information about whether excess packets are tagged with a special IEEE 802.1p tag value.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

q e disp

- 3900
- 9300

**qos excessTagging
enable****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Tags or retags excess packets with a special 802.1p tag value. This special value refers to any packets that are marked as excess that you want to tag.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q e e

Important Considerations

- Excess tagging is disabled by default.
- When you enter this command, you are prompted to enter an IEEE 802.1p tag value for excess packets in the range of 0 through 7, with 0 as the default. For example, if you specify 1, excess packets become background traffic.

Options

Prompt	Description	Possible Values	[Default]
IEEE 802.1p tag value	Tag value to use to tag or retag excess packets	0 – 7	0

**qos excessTagging
disable**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Disables the tagging of excess packets with a special 802.1p tag value.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

q e disa

Important Consideration

- Excess tagging is disabled by default.

- 3900
- 9300

qos statistics interval *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Sets a sampling interval for gathering QoS statistics.

Valid Minimum Abbreviation

q s i

Important Considerations

- The default interval is 5 seconds.
- When you enter this command, the system prompts you to enter the appropriate interval. The existing value appears in brackets.
- A nonzero value shows the byte or packet-count-per-interval period. A zero value shows byte or packet counters.

Options

Prompt	Description	Possible Values	[Default]
Interval	Interval, in seconds, during which you want to gather QoS statistics	0 – 60 seconds	5 (factory default), or current value

qos statistics receive *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- 9400

Displays QoS receive statistics.

Valid Minimum Abbreviation

q s r

- 3900
- 9300

Important Considerations

- The system displays the statistics at the interval that you specified. The default interval is 5 seconds.
- The receive statistics shows the effect of the traffic control services that you configured.

Options

Prompt	Description	Possible Values	[Default]
Bridge ports	Port numbers whose receive statistics you want to display. On the CoreBuilder® 9000, the list of ports includes the front-panel ports and any enabled backplane ports.	<ul style="list-style-type: none"> ■ One or more port numbers ■ all ■ ? (for a list of selectable ports) 	–

Fields in the QoS Receive Statistics Display

Field	Description
droppedPackets	Number of packets that were dropped when they were received
droppedPacketsPeak	Highest number of packets that were dropped on receipt up to this point
flowExcess	Number of flow classifier bytes that are excess
flowExcessPeak	Highest number of flow excess bytes that have been received up to this point
flowReserved	Number of conforming flow classifier bytes that have been received
flowReservedPeak	Highest number of flow classifier bytes that have been received up to this point
nonFlowExcess	Number of nonflow classifier bytes that have been received that are excess
nonFlowExcessPeak	Highest number of nonflow excess bytes that have been received up to this point

Field	Description
nonFlowReserved	Number of conforming non-flow classifier bytes that have been received
nonFlowResvPeak	Peak count: The highest number of conforming nonflow classifier bytes that have been received up to this point
port	If you display statistics for multiple ports, the port number that is associated with the statistics

qos statistics transmit

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays QoS transmit statistics.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q s t

Important Considerations

- The transmit statistics help you track bandwidth utilization and packet loss by physical port and queue (reserved, high, best, and low).
- The RSVP and network control packets go out on the reserved queue.
- When you mark any packet (conforming or excess) as loss eligible, the packet is dropped if the transmit queue for which it is destined is over its threshold. A packet that is marked loss-eligible falls into one of the two highLoss statistic categories:
 - If the transmit queue is not over its threshold, the packet is sent and counted as a highLossSent packet.
 - If the transmit queue is over its threshold, it is dropped and counted as a highLossDropped packet.
- If you do *not* mark a packet as loss-eligible, it falls into one of the three lowLoss statistics.
 - If the queue is not over the threshold, it is counted as a lowLossSent.
 - If the queue is over its threshold, it is counted as lowLossDelayed.
 - If the queue is full, it is counted as lowLossDropped.
- *Loss-eligible packets* are conforming packets that are discarded instead of queued when transmit queues back up beyond a threshold. You can specify whether conforming packets (as well as nonconforming excess packets) are loss-eligible when you define a control. Marking packets loss-eligible is useful to enable intelligent discard of traffic in a congestion situation. When the system is congested, you can decide which traffic can be discarded and mark that traffic as loss eligible.

Options

Prompt	Description	Possible Values	[Default]
Bridge ports	Port numbers of ports for which you want to display transmit statistics. On the CoreBuilder® 9000, the list of ports includes the front-panel ports and any enabled backplane ports.	<ul style="list-style-type: none"> ■ One or more port numbers ■ all ■ ? (for a list of selectable ports) 	–
Queues	Transmit queues (types of service) whose statistics you want to display.	<ul style="list-style-type: none"> ■ reserved ■ high ■ best ■ low ■ all ■ ? (for a list of selectable values) 	–

Fields in the QoS Transmit Statistics Display

Field	Description
highLossDropped	Number of loss-eligible packets that were discarded and were over the threshold
highLossDroppedPeak	Current highest count of loss-eligible packets that were discarded and were over the threshold
highLossSent	Number of loss-eligible packets that were sent and were under the threshold (at low latency)
highLossSentPeak	Current highest count of loss-eligible packets that were sent and were under the threshold
lowLossDelayed	Number of non-loss-eligible packets that were sent and over the threshold (that is, the transmit queues were backing up but not overflowing)
lowLossDelayedPeak	Current highest count of non-loss-eligible packets that were sent and were over the threshold
lowLossDropped	Number of packets that were discarded because they exceeded the length of the transmit queue
lowLossDroppedPeak	Current highest count of packets that were discarded because they exceeded the length of the transmit queue
lowLossSent	Number of non-loss-eligible packets that were sent and were under the threshold (at low latency)


Field	Description
lowLossSentPeak	Current highest count of non-loss-eligible packets that were sent and were under the threshold
port	Port number that is associated with the statistics
queue	Queue that is associated with the statistics



MONITORING

Chapter 23 **Event Log**

Chapter 24 **Roving Analysis**



EVENT LOG

This chapter provides guidelines and other key information about how to administer event logs in your system, including the following tasks:

- Display the event log configuration
- Configure the output devices
- Configure the services

Use event logging to capture different types of log messages from various services (applications) and send them to the Administration Console. The log messages display real-time information about the state of the system or a specific service, and can help you diagnose site-specific problems.



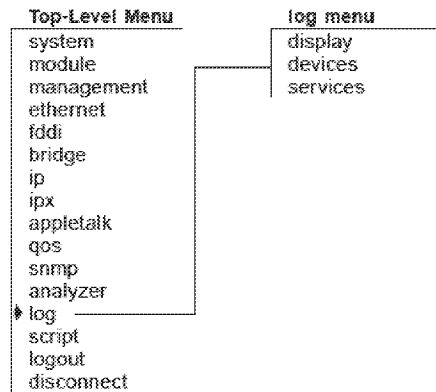
On CoreBuilder® 9000 systems, event logging is controlled entirely through the Enterprise Management Engine (EME), not through the Administration Consoles of individual modules as described here. See the CoreBuilder 9000 Enterprise Management Engine User Guide for information on how to keep logs of switch events.



For more information about implementing event logging on your network, see the CoreBuilder 3500 Implementation Guide.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



✓ 3500
9000
9400

3900
9300

log display Displays the current log settings.

Valid Minimum Abbreviation

log di

Important Consideration

- The CoreBuilder 3500 by default enables logging to the serial port session and disables logging to any Telnet or modem session. However, you can toggle the current logging state on the CoreBuilder 3500 from serial port to Telnet or modem by entering Ctrl+L.

Fields in the Log Display

Field	Description
consoleOut	Administration Console output device. You can enable or disable the Console to display event log messages for each severity level.
Logging message	Whether logging to this console session is enabled or disabled.
Supported Event Log Services	
AppleTalk	Appletalk log service. Enabled or disabled for each severity level.
IPX	IPX log service. Enabled or disabled for each severity level.
System	System log service. Enabled or disabled for each severity level.
Severity Levels	
Config	Configuration changes.
Error	Application-specific error. Default: enabled
Info	Severity level of changes in the state of the system that are not caused by events at any other severity level
Warning	Nonfatal problem. Default: enabled

log devices Configures severity levels for event logging on the Administration Console.

✓ 3500
9000
9400

Valid Minimum Abbreviation

log de

Important Considerations

- You can set the console to log events for one or more of the four severity levels.
- To specify multiple severity levels, separate the levels with a comma (for example, `warning, config`).

Options

Prompt	Description	Possible Values	[Default]
Levels for console	Event logging severity level for console output	<ul style="list-style-type: none"> ■ error ■ warning ■ config ■ info ■ all ■ ? (for a list of valid severity levels) 	–
Selected levels	Whether selected event logging for console output is enabled or disabled	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y

Log Devices Examples (3500)

Select menu option (log): `devices`

Select levels for console (`error,warning,config,info|all|?`): `?`

Selectable values

`error, warning, config, info`

Select levels for console (`error,warning,config,info|all|?`): `all`

Enable the selected levels (`n,y`) [`y`]: `y`

To disable the config and info severity levels:

```
Select menu option (log): devices  
Select levels for console (error,warning,config,info|all|?): config,info  
Enable the selected levels (n,y) [y]: n
```

The display now indicates that the error and warning severity levels remain enabled and the config and info levels are disabled.

- ✓ **3500**
9000
9400
- log services** Enables the logging of messages that pertain to the following services:
- System level
 - AppleTalk
 - IPX

3900
9300

Valid Minimum Abbreviation

log s

Important Considerations

- For a specific service or all services, you can configure up to four severity levels.
- Use a comma to separate multiple service names and severity levels (for example, `system, appletalk and error, warning`).

Options

Prompt	Description	Possible Values	[Default]
Services	Services to configure	<ul style="list-style-type: none"> ■ system ■ ipx ■ appletalk ■ all ■ ? (for a list of valid services to configure) 	–
Levels	Severity levels to enable	<ul style="list-style-type: none"> ■ error ■ warning ■ config ■ info ■ all ■ ? (for a list of valid severity levels to enable) 	–
Selected services/levels	Whether the selected services and severity levels are enabled or disabled	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y

Log Services Examples

To enable all severity levels for the AppleTalk service:

```
Select menu option (log): services
```

```
Select services (system,ipx,appletalk|all|?): ?
```

Selectable values

```
system,ipx,appletalk
```

```
Select services (system,ipx,appletalk|all|?): appletalk
```

```
Select levels (error,warning,config,info|all|?): all
```

```
Enable the selected services/levels (n,y) [y]: y
```

To show that all severity levels are enabled for the AppleTalk service, enter `log display`

To disable the warning and info severity levels for the AppleTalk service, follow this example:

```
Select menu option (log): services
```

```
Select services (system,ipx,appletalk|all|?): appletalk
```

```
Select levels (error,warning,config,info|all|?): warning,info
```

```
Enable the selected services/levels (n,y) [y]: n
```

To show that the AppleTalk service is associated with only the error and config severity levels, enter `log display`

ROVING ANALYSIS

This chapter provides guidelines and other key information about how to set up roving analysis in your system, including the following tasks:

- Display roving analysis configuration
- Add and remove analyzer
- Start and stop monitoring

Roving analysis is the mirroring of traffic on one port to another port of the same media type.

- The port being monitored is called the *monitor port*.
- The port that receives the mirrored traffic is called the *analyzer port*.

The analyzer port typically has a network analyzer or RMON *probe* attached through which you can watch the network traffic.

Use roving analysis to monitor Fast Ethernet, Gigabit Ethernet, or Fiber Distributed Data Interface (FDDI) port traffic for network management and troubleshooting purposes. You use the Administration Console to choose any network segment that is attached to a system and monitor its activity.

You can monitor a designated roving analysis port to:

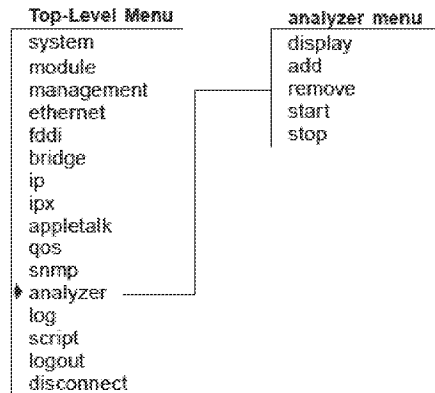
- Analyze traffic loads on each segment so that you can continually optimize your network loads by moving network segments
- Troubleshoot network problems (for example, to find out why a particular segment has so much traffic)



For more information about implementing roving analysis on your network, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



analyzer display

Displays the roving analysis configuration, showing which ports are designated as analyzer ports and which bridge ports are currently being monitored.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

an d

- ✓ 3900
- ✓ 9300

Fields in the Analyzer Display

Field	Description
Ports configured as analyzer ports	List of analyzer ports on the system, including the port number and MAC address. These are the ports that can accept traffic that is mirrored from a monitored port. Analyzer ports are typically connected to a network analyzer or probe. There may be multiple analyzer ports defined on the switch.
Port	Analyzer port number
Type	Media type and Port Speed (FDDI, Fast Ethernet, or Gigabit Ethernet)
Address	MAC address of the analyzer port
Ports being monitored	List of ports that the system is monitoring. Includes the MAC address of the analyzer port to which the monitored port traffic will be forwarded.
Port	Monitored port number
Type	Media type and Port Speed (FDDI, Fast Ethernet, or Gigabit Ethernet)
Analyzer Address	MAC address of the analyzer port to which the monitored port traffic will be forwarded and to which your network analyzer or probe is attached. There may be multiple analyzer ports defined on the switch.

Analyzer Display Example (3500)

Select menu option (analyzer): display

Ports configured as analyzer ports:

Port	Type	Address
8	Fast Ethernet	00-80-3e-2b-42-08

Ports being monitored:

Port	Type	Analyzer Address
12	Fast Ethernet	00-80-3e-2b-42-08

analyzer add Defines a bridge port to serve as a dedicated analyzer port.

✓ 3500
 ✓ 9000
 ✓ 9400

Valid Minimum Abbreviation

an a

Important Considerations

✓ 3900
 ✓ 9300

- On CoreBuilder® 3500 and CoreBuilder 9000 systems, you can connect as many as 16 network analyzers to a system. On other platforms, you can connect one network analyzer. For more accurate analysis, attach the analyzer to a dedicated port instead of through a repeater.
- After a port is selected to serve as an analyzer port, it cannot receive or transmit any other data. Instead, it receives only the data from the ports to be monitored. If you have enabled the Spanning Tree Protocol (STP) on the port, STP is automatically disabled.
- If the physical port configuration changes in the system (that is, if you remove or rearrange modules), the MAC address of the analyzer port remains fixed. If you replace the module with the analyzer port with a module of a different media type, the roving analysis port (RAP) configuration for that port is cleared.
- When you configure a port that is part of a virtual LAN (VLAN) as an analyzer port, a warning is displayed because adding the port removes the port from all VLANs. When the port is restored (when you remove the analyzer port), it becomes a member of the default VLAN.
- If the probe is attached to a 10 Mbps Ethernet analyzer port and the roving analysis port (RAP) is monitoring a 100 Mbps Ethernet port with a sustained traffic rate greater than 10 Mbps, the analyzer may not see all of the frames.
- After you enter a bridge port number, the system displays the MAC address of the analyzer port. Record this information for setting up the port that you want to monitor.
- On the CoreBuilder 9000, the port to which the analyzer is attached and the port you wish to monitor must be on the same blade.
- Trunked ports and resilient link ports can not be configured as analyzer ports.

Options

Prompt	Description	Possible Values	[Default]
Bridge port	Number of the bridge port to which you want to attach the analyzer n varies by platform. Only valid port number choices are displayed.	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of available bridge ports) 	–

Analyzer Add Example (9000 1000BASE-SX module)

```

CB9000@slot 3.1 [9-GEN-SX-L2] (): analyzer add
Select bridge port {1-9|?}: 9
Warning: Port being removed from Vlan: Default
Analyzer port address is 00-20-9c-0d-e1-2a

```

analyzer remove Restores the port to be a regular bridge port. Restores the Spanning Tree state to its state before the port was configured as an analyzer port.

✓ 3500
 ✓ 9000
 ✓ 9400

Valid Minimum Abbreviation

an r

✓ 3900
 ✓ 9300

Important Considerations

- Use this command when you no longer need the bridge port for the analyzer.
- The analyzer port can not be removed if it still has monitor ports.
- The port becomes a member of the default virtual LAN (VLAN) when it is restored (when you remove it as an analyzer port).
- The port will not be automatically restored to any VLAN it might have been a member of before it was configured as an analyzer port — you must do this yourself.

Options

Prompt	Description	Possible Values	[Default]
Bridge port	Number of the bridge port to which the analyzer is attached n varies by platform. Only active analyzer port numbers are displayed.	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of available bridge ports) 	–

Analyzer Remove Example (3500)

```
Select menu option (analyzer): remove
Select bridge port {2,7|?}: 7
```

analyzer start Starts port monitoring activity on the selected bridge port.

✓ 3500
 ✓ 9000
 ✓ 9400

Valid Minimum Abbreviation

an sta

Important Considerations

- ✓ 3900
- ✓ 9300
- You must already have an analyzer port configured. First designate a bridge port to serve as the analyzer port and connect the analyzer to that port. See “analyzer add” earlier in this chapter for details.
- On the CoreBuilder 9000, the analyzer port and the monitor port must be on the same module.
- The MAC address of the analyzer port is displayed when you configure that port, and when you display the roving analysis configurations on the system to which the analyzer is attached.
- The media type of the analyzer port must match the media type of the port being monitored. Fast Ethernet and Gigabit Ethernet are the same media type.
- You can use a Fast Ethernet (10 Mbps) port to monitor a Gigabit Ethernet (100 Mbps) port, but a warning message will be printed. If the sustained traffic load is greater than 10 Mbps, the analyzer on the slower port may not see all the frames on the faster port.
- When you successfully configure a bridge port to be monitored, all the data that the monitored port receives and transmits is copied to the selected analyzer port.
- Once a port is selected to serve as a monitor port, the RMON data that it can record is limited to the RMON groups (statistics, history, alarm, event, protocolDir, and probeConfig) that do not require hardware sampling.
- If you replace the module that the monitored port resides on with a module of a different media type, the roving analysis port (RAP) configuration for the monitored port is reset.

Options

Prompt	Description	Possible Values	[Default]
Bridge port	Number of the bridge port to be monitored n varies by platform.	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of available bridge ports) 	–
Target analyzer port address	MAC address of the port to which the analyzer is attached	A valid MAC address of an analyzer port	–

Analyzer Start Example (9000 1000BASE-SX module)

```

CB9000@slot 3.1 [9-GEN-SX-L2] (analyzer): start
Select bridge port {1-8,10-12|?}: 1
Enter the target analyzer port address: 00-20-9c-0d-e1-2a

```

analyzer stop Stops port monitoring activity on the selected bridge port.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

an sto

Important Consideration

- Port data is no longer copied and forwarded to the selected analyzer port from the port that you specify. See “analyzer start” earlier in this chapter for details.


- ✓ 3900
- ✓ 9300

Options

Prompt	Description	Possible Values	[Default]
Bridge port	Number of the bridge port that is being monitored n varies by platform.	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of available bridge ports) 	–


Analyzer Stop Example (3500)

```
Select menu option (analyzer): stop
Select bridge port {3,4|?}: 3
```

REFERENCE

Appendix A Technical Support





TECHNICAL SUPPORT

3Com provides easy access to technical support information through a variety of services. This appendix describes these services.

Information contained in this appendix is correct at time of publication. For the most recent information, 3Com recommends that you access the 3Com Corporation World Wide Web site.

Online Technical Services

3Com offers worldwide product support 24 hours a day, 7 days a week, through the following online systems:

- World Wide Web site
- 3Com Knowledgebase Web Services
- 3Com FTP site
- 3Com Bulletin Board Service (3Com BBS)
- 3Com Facts™ Automated Fax Service

World Wide Web Site

To access the latest networking information on the 3Com Corporation World Wide Web site, enter this URL into your Internet browser:

<http://www.3com.com/>

This service provides access to online support information such as technical documentation and software, as well as support options that range from technical education to maintenance and professional services.

3Com Knowledgebase Web Services

This interactive tool contains technical product information compiled by 3Com expert technical engineers around the globe. Located on the World Wide Web at <http://knowledgebase.3com.com>, this service gives all 3Com customers and partners complementary, round-the-clock access to technical information on most 3Com products.

3Com FTP Site Download drivers, patches, software, and MIBs across the Internet from the 3Com public FTP site. This service is available 24 hours a day, 7 days a week.

To connect to the 3Com FTP site, enter the following information into your FTP client:

- Hostname: `ftp.3com.com`
- Username: `anonymous`
- Password: `<your Internet e-mail address>`



You do not need a user name and password with Web browser software such as Netscape Navigator and Internet Explorer.

3Com Bulletin Board Service The 3Com BBS contains patches, software, and drivers for 3Com products. This service is available through analog modem or digital modem (ISDN) 24 hours a day, 7 days a week.

Access by Analog Modem

To reach the service by modem, set your modem to 8 data bits, no parity, and 1 stop bit. Call the telephone number nearest you:

Country	Data Rate	Telephone Number
Australia	Up to 14,400 bps	61 2 9955 2073
Brazil	Up to 28,800 bps	55 11 5181 9666
France	Up to 14,400 bps	33 1 6986 6954
Germany	Up to 28,800 bps	4989 62732 188
Hong Kong	Up to 14,400 bps	852 2537 5601
Italy	Up to 14,400 bps	39 2 27300680
Japan	Up to 14,400 bps	81 3 5977 7977
Mexico	Up to 28,800 bps	52 5 520 7835
P.R. of China	Up to 14,400 bps	86 10 684 92351
Taiwan, R.O.C.	Up to 14,400 bps	886 2 377 5840
U.K.	Up to 28,800 bps	44 1442 438278
U.S.A.	Up to 53,333 bps	1 847 262 6000

Access by Digital Modem

ISDN users can dial in to the 3Com BBS using a digital modem for fast access up to 64 Kbps. To access the 3Com BBS using ISDN, call the following number:

1 847 262 6000

3Com Facts Automated Fax Service

The 3Com Facts automated fax service provides technical articles, diagrams, and troubleshooting instructions on 3Com products 24 hours a day, 7 days a week.

Call 3Com Facts using your Touch-Tone telephone:

1 408 727 7021

Support from Your Network Supplier

If you require additional assistance, contact your network supplier. Many suppliers are authorized 3Com service partners who are qualified to provide a variety of services, including network planning, installation, hardware maintenance, application training, and support services.

When you contact your network supplier for assistance, have the following information ready:

- Product model name, part number, and serial number
- A list of system hardware and software, including revision levels
- Diagnostic error messages
- Details about recent configuration changes, if applicable

If you are unable to contact your network supplier, see the following section on how to contact 3Com.

Support from 3Com

If you are unable to obtain assistance from the 3Com online technical resources or from your network supplier, 3Com offers technical telephone support services. To find out more about your support options, call the 3Com technical telephone support phone number at the location nearest you.

When you contact 3Com for assistance, have the following information ready:

- Product model name, part number, and serial number
- A list of system hardware and software, including revision levels
- Diagnostic error messages
- Details about recent configuration changes, if applicable

Here is a list of worldwide technical telephone support numbers:

Country	Telephone Number	Country	Telephone Number
Asia, Pacific Rim			
Australia	1 800 678 515	P.R. of China	10800 61 00137 or 021 6350 1590
Hong Kong	800 933 486	Singapore	800 6161 463
India	+61 2 9937 5085	S. Korea	From anywhere in S. Korea: 00798 611 2230 From Seoul: (0)2 3455 6455
Indonesia	001 800 61 009	Taiwan, R.O.C.	0080 611 261
Japan	0031 61 6439	Thailand	001 800 611 2000
Malaysia	1800 801 777		
New Zealand	0800 446 398		
Pakistan	+61 2 9937 5085		
Philippines	1235 61 266 2602		
Europe			
From anywhere in Europe, call:	+31 (0)30 6029900 phone		
	+31 (0)30 6029999 fax		
Europe, South Africa, and Middle East			
From the following countries, you may use the toll-free numbers:			
Austria	0800 297468	Netherlands	0800 0227788
Belgium	0800 71429	Norway	800 11376
Denmark	800 17309	Poland	00800 3111206
Finland	0800 113153	Portugal	0800 831416
France	0800 917959	South Africa	0800 995014
Germany	0800 1821502	Spain	900 983125
Hungary	00800 12813	Sweden	020 795482
Ireland	1800 553117	Switzerland	0800 55 3072
Israel	1800 9453794	U.K.	0800 966197
Italy	1678 79489		
Latin America			
Argentina	AT&T +800 666 5065	Mexico	01 800 CARE (01 800 2273)
Brazil	0800 13 3266	Peru	AT&T +800 666 5065
Chile	1230 020 0645	Puerto Rico	800 666 5065
Colombia	98012 2127	Venezuela	AT&T +800 666 5065
North America			
	1 800 NET 3Com (1 800 638 3266)		
	Enterprise Customers: 1 800 876-3266		

Returning Products for Repair

Before you send a product directly to 3Com for repair, you must first obtain an authorization number. Products sent to 3Com without authorization numbers will be returned to the sender unopened, at the sender's expense.

To obtain an authorization number, call or fax:

Country	Telephone Number	Fax Number
Asia, Pacific Rim	+ 65 543 6500	+ 65 543 6348
Europe, South Africa, and Middle East	+ 31 30 6029900	+ 31 30 6029999
Latin America	1 408 326 2927	1 408 326 3355

From the following countries, you may call the toll-free numbers; select option 2 and then option 2:

Austria	0800 297468	
Belgium	0800 71429	
Denmark	800 17309	
Finland	0800 113153	
France	0800 917959	
Germany	0800 1821502	
Hungary	00800 12813	
Ireland	1800553117	
Israel	1800 9453794	
Italy	1678 79489	
Netherlands	0800 0227788	
Norway	800 11376	
Poland	00800 3111206	
Portugal	0800 831416	
South Africa	0800 995014	
Spain	900 983125	
Sweden	020 795482	
Switzerland	0800 55 3072	
U.K.	0800 966197	
U.S.A. and Canada	1 800 NET 3Com (1 800 638 3266)	1 408 326 7120 (not toll-free)
	Enterprise Customers: 1 800 876 3266	

INDEX

Symbols

? character 473, 478

Numbers

3C number 69, 129
3Com bulletin board service (3Com BBS) 768
3Com Knowledgebase Web Services 767
3Com URL 767
3ComFacts 769
802.3_RAW packets 256

A

AARP (AppleTalk Address Resolution Protocol) 674 to 676
access levels 37, 38
 and passwords 75
addModify (snmp trap) 195
address group
 adding port addresses 391
address threshold 252
address/port patterns for QoS classifiers 696, 701, 703
addresses
 adding static 294
 for SNMP trap reporting 194
addressThresholdEvent 257
Administration Console 29 to 42
 password access 75
administration console
 of an ATM switch 30
advancedPing 179, 475, 476
advancedTraceRoute 184
 packet size 480
 ttl option 480
 wait option 480
advertise RIP mode 451, 453
AEP (AppleTalk Echo Protocol) 682
aggregate rate limit 710, 712
 for flow classifiers 715
aggregated links 299
aging time 258

allClosed mode for VLANs
 and Ignore STP mode 365
 displaying 338
 selecting 364
allOpen mode for VLANs
 displaying 338
 selecting 364
analyzer port
 MAC address 758
anchor ports
 rate limits affecting 716
 trunking 313
AppleTalk
 AARP (AppleTalk Address Resolution Protocol) 674 to 676
 AEP (AppleTalk Echo Protocol) 682
 checksums 680
 DDP statistics 683
 forwarding 679
 interfaces 663 to 670
 NBP statistics 686
 ping 682
 removing interfaces 669
 routes 672, 673
 RTMP statistics 684
 source socket verification 681
 ZIP statistics 685
 zones 677, 678
applying controls to classifiers 713
areas 531 to 537
ARP (Address Resolution Protocol)
 cache 171
 deleting cache entries 170
 deleting dynamic cache entries 172
 displaying cache 168
 flushing all entries 171
 flushing dynamic entries 172
 remove 170
 static cache entry 169
ARP cache 428 to 433
ASCII-based editor
 and scripts 124
ATM switch 30
autonegotiation, Ethernet ports 212

-
- B**
- backplane ports, interface module 31
 - backup
 - saving NV data 107
 - bandwidth, QoS
 - displaying 733
 - modifying 734
 - bandwidth, RSVP 725, 728
 - baseline, setting current 133
 - baud rate
 - serial port 94, 97
 - baud setting 95
 - best service level 711
 - blocking, ignoring STP 365
 - BOOTP (Boot Protocol)
 - as UDP service 442
 - hop count 442
 - relay threshold 446
 - bridge ports
 - adding MAC addresses 294
 - defining VLANs 345, 352
 - deleting VLANs 363
 - listing MAC addresses 293
 - modifying VLANs 355, 360
 - VLAN summary 339, 342
 - bridge-wide parameters, allOpen or allClosed VLAN mode 364
 - bulletin board service 768
 - burst size, QoS control 712
 - burst, advancedPing option 476
-
- C**
- cast types for QoS classifiers 695, 718
 - changing VLANs 355, 360
 - channels
 - management and data 31
 - chassis
 - management architecture 33
 - power management 33
 - checksums, AppleTalk 680
 - Class of Service 267
 - Class of Service (CoS) 267
 - classifiers, QoS
 - applying controls to 713
 - default 694
 - defining 694
 - displaying detail information 692
 - displaying summary information 691
 - guidelines for using 689
 - modifying 701
 - parameters for defining 695, 718
 - predefined flow and nonflow 691
 - removing 706
 - specifying address/port patterns 696, 701, 703
 - specifying IP addresses 695, 702
 - command strings
 - entering abbreviated 41
 - entering values 41
 - quick 34
 - commands 150
 - system menu
 - for baselining statistics 90
 - for managing NV data 108
 - community strings 192
 - configuration tasks 34
 - conforming packets
 - service levels 711
 - console access 73
 - control packets 733, 734
 - controls, QoS
 - associating with classifiers 713
 - burst size 712
 - defining 710
 - displaying detail information 708
 - displaying summary information 707
 - modifying 718, 721
 - names 711
 - parameters for defining 710
 - removing 724
 - service levels 711
 - specifying IEEE 802.1p tags 712
 - conventions
 - notice icons 23
 - text 24
 - CoreBuilder 3500 system
 - and network monitoring 755
 - CoreBuilder 9000
 - management features 33
 - system management overview 30
 - cost
 - IP RIP mode 456
 - Spanning Tree settings 254
-
- D**
- DAS (dual attach station) pairs
 - trunks and 307
 - data channels
 - management 31
 - date
 - displaying 137
 - DDP (Datagram Delivery Protocol) 683
 - defaults
 - control service level (best) 715

- IP RIP mode (learn) 453
- OSPF route metric 538 to 540
- QoS classifier 694
- route for IP 421
- screen height 76
- Spanning Tree Protocol 261
- ttl value for advancedTraceRoute 480
- ttl value for traceRoute 182
- UDP port number for advancedTraceRoute 480
- UDP port number for traceRoute 182
- defining
 - QoS controls 710, 712, 715
 - VLANs 310, 331
- deleting
 - links 336
 - trunks 318
 - VLANs 363
- designated root 253
- destination address
 - for SNMP trap reporting 194
- destination IP address for QoS classifiers 702
- destination IP address masks 702
- detail
 - trunks 329
- detail information
 - trunks 305
 - VLANs 341
- details, AppleTalk interface 664
- Diagnostics status 69
- disabled RIP mode 451, 453
- disabling
 - excess packet tagging 737
 - RSVP 732
- displaying
 - QoS bandwidth 733
 - QoS classifier detail 692
 - QoS classifier summary 691
 - QoS control detail information 708
 - QoS excess packet tagging 735
 - QoS summary information 691, 707
 - RSVP detail information 729
 - summary RSVP information 728
- displaying TCMP state 301, 304, 328
- DNS (Domain Name System) servers 436 to 440
- documentation
 - comments 25
- drop service level 711
- duplex mode, Ethernet ports 212, 213
- DVMRP (Distance-Vector Multicast Routing Protocol) 507
- dynamic versus static VLAN origin 345, 352

E

- edge policing option 725, 728
- editor for scripts
 - EMACS 124
 - vi 124
- EME (Enterprise Management Engine)
 - console 30
 - overview 33
- enabled RIP mode 451, 453
- enabling
 - excess packet tagging 736
 - RSVP 730
- enabling and disabling Ethernet ports 220
- errors
 - routing interface 406
- Ethernet
 - address
 - and restoring NV data 110
 - and roving analysis 757
 - autonegotiation 212
 - enabling and disabling ports 220
 - fragmenting packets 255
 - menu options 203
 - PACE Access 217
 - PACE Interactive Access 218
 - port duplex mode 212, 213
 - port flow control 215
 - port labels 219
 - port monitoring 221, 222
 - port numbering 204, 207
 - port speed 212, 213
 - port state 220
 - statistics 204, 208
- event log 747, 752
- examples
 - defining QoS classifiers 699
 - defining QoS controls 717
 - defining VLANs (Layer 2 devices) 354
 - defining VLANs (Layer 3 devices) 350, 351, 358, 359
 - modifying QoS classifiers 705
 - modifying QoS controls 722
 - modifying VLANs (Layer 2 devices) 362
 - modifying VLANs (Layer 3 devices) 359, 362
 - of a script 125
 - removing VLANs 363
 - setting ignore STP mode 365
 - trunk changes 316
 - trunk definitions 311
- excess packet tagging, QoS 735, 737
- excess packets
 - treatment with RSVP 728, 730
- extended diagnostics version number 69

F

- fax service (3ComFacts) 769
- FDDI (Fiber Distributed Data Interface)
 - fragmenting packets 255
 - port label 241
- FDDI MAC
 - condition report 237
 - LLC Service, enabling 239
- FDDI station
 - and SRFs 224, 228
- FDDI_Snap packets 256
- feedback on documentation 25
- File Transfer Protocol (FTP) 87, 89, 107
- filter id 371
- filters for QoS flow classifiers
 - defining 696
 - modifying 701
- flow classifiers
 - cast types 695, 701, 702
 - defining 694
 - predefined 691
 - protocol types 695, 701, 702
 - removing 706
 - using aggregate rate limit 715
- flow control
 - defining for Gigabit trunk 310, 315
 - displaying for trunks 305, 329
- flow control, Ethernet ports 215
- flush
 - for management ip routes 164
 - snmp trap 197
- flushing
 - learned IP routes 424
 - SNMP trap addresses 197
- forwarding
 - AppleTalk 679
 - IPX 629

G

- gateway IP address 421
- Gigabit Ethernet
 - trunks 318
- guidelines
 - for using QoS 689
- GVRP (GARP VLAN Registration Protocol)
 - displaying status 338
 - using 345, 352

H

- hardware revision numbers 69
- high service level 711
- hop count 442

I

- ICMP statistics 187, 483
- ID, VLAN 341
- IEEE 802.1p priority tagging
 - for excess packets 735, 736
 - for nonflow classifiers 697, 703, 705
 - for QoS controls 712
- IEEE 802.1Q tagging 349
- IGMP (Internet Group Management Protocol)
 - query mode 524, 525
 - snooping mode 524, 525
- IGMP snooping, Layer 2 devices 270
- IGMP snooping, Layer 3 devices 524
- ignore STP mode
 - selecting 365
- ignoring blocking for VLANs 365
- in-band-management 149
- index 356
 - VLAN interface 360, 395
- interface module
 - backplane ports 31
- interfaces
 - IP 449
 - OSPF 541 to 556
- interfaces, AppleTalk
 - define 665
 - detail display 664
 - modify 667
- interval, QoS statistics 738
- IP (Internet Protocol)
 - address masks for QoS classifiers 695
 - addresses 395, 421
 - loading software 89
 - QoS classifiers 702
 - addresses and restoring NV data 110
 - advancedPing 475
 - advancedTraceRoute 480
 - ARP cache 428 to 433
 - defining routes 422
 - DNS 436
 - enabling or disabling routing 450
 - interfaces
 - displaying 608, 610
 - removing 613, 649
 - statistics 418
 - summary information 398

- overlapped interfaces 447 to 449
 - overview 149
 - ping functions 473, 478
 - RIP mode 451
 - routes 450
 - statistics 482
 - ICMP 483
 - UDP 483
 - traceRoute functions 478
 - UDP Helper 442, 447
 - IP multicast
 - cache 518
 - DVMRP metric 507, 510
 - hop count 517
 - IGMP 524, 525
 - prune messages 518, 519
 - routing table 517
 - TTL threshold 510
 - tunnels 511, 513, 517, 519, 527
 - IP multicast filtering
 - IGMP snooping 524, 525
 - IP multicast routing
 - DVMRP 507
 - IGMP 525
 - routeDisplay 517
 - IP protocol types
 - modifying 704
 - IP routes
 - flushing 424
 - interface status 421
 - IP routing
 - enabling or disabling 450
 - IPX
 - forwarding
 - enabling or disabling 629
 - statistics 653
 - interfaces
 - statistics 655
 - RIP mode
 - setting 630
 - statistics 651
 - triggered updates 631
 - RIP policy
 - define 633
 - summary 632, 637
 - routes
 - defining static 618
 - flushing learned routes 621
 - removing 620
 - SAP (Service Advertisement Protocol) mode
 - statistics 652
 - triggered updates 639
 - SAP mode 638
 - SAP policy
 - define 642
 - detail 641
 - modify 645
 - remove 648
 - summary 640
 - static servers 622, 624
-
- ## K
- KBytes/sec rate limit 712
-
- ## L
- labels, Ethernet ports 219
 - LANs
 - virtual 337
 - Layer 2 devices
 - defining VLANs 352
 - modifying VLANs 360
 - Layer 3 addresses
 - for VLANs 347
 - modifying 355
 - Layer 3 devices
 - defining VLANs 346
 - modifying VLANs 355
 - learn RIP mode 451, 453
 - learned routes, IP 424
 - learning state 264
 - LER (Link Error Rate)
 - alarm value 242
 - lerCutoff
 - and lerAlarm value 243
 - levels, service 711
 - limits
 - for QoS classifiers 694
 - QoS rate 710, 712, 715
 - link aggregation 299
 - link state database, OSPF 557 to 563
 - links
 - removing 336
 - resources 336
 - listening state 264
 - LLC (Logical Link Control)
 - service description 239
 - LMA (Local Management Application), ATM Switch 31
 - log, event 747
 - logout 126
 - low service level 711

-
- M**
- MAC (Media Access Control) addresses
 - adding 294
 - displaying 293
 - MAC type for trunk 310, 315
 - management
 - and naming the system 101, 136
 - configuring system access 190
 - displaying detailed information 153
 - displaying summary information 151
 - SNMP community strings 192
 - Transcend Network Control Services 30
 - Web Management applications 30
 - management data channels 31
 - management ip
 - advancedPing 179
 - advancedTraceRoute 184
 - displaying statistics 186
 - ping 177
 - statistics 176
 - tracing a route destination 182
 - management ip arp
 - defining a static cache entry 169
 - displaying cache 168
 - flushing all entries from cache 171
 - flushing dynamic entries 172
 - removing cache entries 170
 - management ip interface
 - defining the IP address 157
 - displaying summary information 156
 - modifying 158
 - removing 159
 - management ip rip
 - displaying RIP information 173
 - management ip route
 - default 165
 - defining a static route 162
 - deleting default 166
 - displaying the routing table 160
 - finding in table 167
 - flushing learned routes 164
 - noDefault 166
 - removing an existing route 163
 - searching the routing table 167
 - masks
 - source and destination IP address 695, 702
 - subnet 395
 - maximum per-reservation bandwidth 730
 - maximum total reservable bandwidth 730
 - memory partition, OSPF 569, 570
 - memory size 69
 - menu structure 150
 - menus
 - and command strings 40
 - entering abbreviated command strings 41
 - entering values 41
 - navigating 42
 - selecting options 40
 - MIBs 768
 - MLAN channel 31
 - mode, operating
 - defining for Ethernet trunk 310, 315
 - displaying for trunks 305, 329
 - modem
 - external, configuring 99, 100
 - modes, VLAN
 - definition 340, 343
 - displaying 338
 - selecting allOpen or allClosed 364
 - selecting Ignore STP 365
 - modifying
 - QoS bandwidth 734
 - QoS classifiers 701
 - QoS controls 718
 - VLANs (Layer 2 devices) 360
 - VLANs (Layer 3 devices) 355, 358
 - module
 - diagnostic messages 129
 - displaying date 137
 - module status information 69
 - monitoring
 - ports, Ethernet 221, 222
 - MultiPoint Link Aggregation (MPLA) 321
 - mode 324
 - Peer Switch Interface State 322
-
- N**
- name server, DNS 436 to 440
 - names
 - for QoS classifiers 695, 718
 - for QoS controls 710
 - trunk 316
 - VLAN 349
 - navigating menus 42
 - NBP (Name Binding Protocol) 686
 - neighbor notification
 - and LLC Service 239
 - neighbors, OSPF 564 to 566
 - network supplier support 769
 - network troubleshooting 755
 - none, for rate limit 710, 712, 715
 - nonflow classifiers
 - cast types 695, 701, 702
 - defining 694

- predefined 691
- protocol types 695, 701, 702
- removing 706
- specifying IEEE 802.1p tags 697, 703, 705
- numbering
 - ports, Ethernet 204, 207
- numbers
 - for QoS classifiers 695, 718
 - for QoS controls 710
- NV data
 - restoring 110

O

- online technical services 767
- origin, VLAN 340, 343
- OSPF (Open Shortest Path First)
 - areas 531 to 537
 - default route metric 538 to 540
 - interfaces 541 to 556
 - link state database 557 to 563
 - memory partition 569, 570
 - neighbors 564 to 566
 - router ID 567
 - routing policies 590 to 602
 - soft restarts 570
 - statistics 603
 - stub default metrics 571 to 573
 - virtual links 574 to 589
- out-of-band
 - management 149
- overlapped IP interfaces 447 to 449

P

- PACE Access, Ethernet 217
- PACE Interactive Access, Ethernet 218
- packet filter
 - displaying contents 372, 373, 374, 376, 377, 379, 382, 384
 - filter id 371
 - processing paths 382
- packet size
 - advancedPing 475
 - advancedTraceRoute 480
- packets
 - tagging of excess 736, 737
- password
 - access levels 35
 - configuring 75
 - IP RIP-2 interface 459
- percentage rate limit 712
- per-reservation bandwidth 725, 728

- ping 177
 - advanced ping example 181
 - example 178
- ping command
 - possible responses 473
- pings, AppleTalk 682
- policing options, RSVP 725, 728
- policy
 - IPX RIP
 - define 633
 - modify 635
 - summary 632, 637
 - IPX SAP
 - define 642
 - detail 641
 - modify 645
 - remove 648
 - summary 640
- policy-based services 689
- port
 - label 241
- port group
 - adding ports 391
- port number
 - setting the traceRoute 182, 480
- port ranges for QoS flow classifiers 696, 701, 703
- port speed 95
 - terminal port, setting the 93, 96
- port state, Ethernet 220
- ports
 - autonegotiation, Ethernet 212
 - defining for VLANs 346
 - defining in trunks 310, 315
 - duplex mode, Ethernet 212, 213
 - enabling and disabling, Ethernet 220
 - flow control, Ethernet 215
 - labels, Ethernet 219
 - maximum number in group 391
 - monitoring, Ethernet 221, 222
 - numbering, Ethernet 204, 207
 - PACE Access, Ethernet 217
 - PACE Interactive Access, Ethernet 218
 - receive ports for controls 712, 716
 - speed, Ethernet 212, 213
 - speed, setting 95
 - state, Ethernet 220
 - statistics, Ethernet 204, 208
 - tagging 349
- predefined QoS classifiers 691
- priority tags
 - excess packets 736
 - nonflow classifiers 697, 703, 705
 - QoS controls 712

prioritization 267
 probe
 RMON 755
 procedures
 defining controls 715
 defining flow classifiers 697
 defining nonflow classifiers 700
 defining RSVP 731
 defining VLANs (Layer 2 devices) 354
 defining VLANs (Layer 3 devices) 349
 modifying VLANs (Layer 2 devices) 361
 modifying VLANs (Layer 3 devices) 358
 protocol types
 for QoS classifiers 695, 718
 modifying for VLANs 355
 modifying QoS classifier 704
 selecting for VLANs 345, 347, 352
 prune messages
 IP multicast 519

Q

QoS (Quality of Service) bandwidth
 displaying 733
 modifying 734
 QoS (Quality of Service) classifiers
 defining 694
 displaying detail information 692
 displaying summary information 691
 example of defining 699
 example of modifying 705
 guidelines for using 689
 modifying 701
 removing 706
 QoS (Quality of Service) controls
 applying to classifiers 713
 defining 710
 displaying detail information 708
 displaying summary information 707
 example of defining 717
 example of modifying 722
 modifying 718, 721
 removing 724
 service levels 711
 specifying rate limits 710, 712, 715
 QoS (Quality of Service) excess packet tagging
 disabling 737
 displaying 735
 enabling 736
 QoS (Quality of Service) statistics
 interval 738
 receive 739
 transmit 741

quiet
 advancedPing option 475

R

rate limits, QoS control 710, 712, 715
 modifying one or more 718
 using with trunks 716
 reboots
 trunks and 307
 receive ports
 rate limit 712
 specifying for trunks 716
 receive statistics, QoS 739
 receivePort rate limit 710, 712, 715
 relay threshold
 BOOTP 446
 remote access 73
 removing
 IP interfaces 612, 613, 649
 links 336
 QoS classifiers 706
 QoS controls 724
 trunks 318
 VLANs 363
 reserved packets 733, 734
 returning products for repair 771
 RIP (Routing Information Protocol)
 display 173
 management statistics 176
 mode example 175
 modes 173, 174
 RIP mode
 IP
 interface information 451
 IPX
 setting 630
 statistics 651
 triggered updates 631
 RIP policy
 define
 IPX 633
 modify 635
 summary
 IPX 632, 637
 RIP-2 password 459
 rlogin
 and rebooting the system 123
 router ID, OSPF 567
 routes
 adding default 165
 AppleTalk 672, 673
 defining static IP 422

- deleting default 166
- finding in table 167
- flushing from the routing table 164
- IPX
 - displaying in routing table 615
 - flushing all learned 621
 - removing 620
 - SAPadvertising 613
 - types of 421
- routing policies, OSPF 590 to 602
- routing analysis
 - and Spanning Tree 758
- RSVP (Resource Reservation Protocol)
 - definition of 689
 - disabling 732
 - displaying detail information 729
 - displaying summary information 728, 729
 - enabling 730
 - policing options 725, 728
 - procedure for defining 731
 - session information 729
 - treatment of excess packets 728, 730
- RTMP (Routing Table Maintenance Protocol) 684

S

- SAP (Service Advertisement Protocol) mode
 - statistics 652
 - triggered updates 639
- SAP mode
 - IPX 638
- SAP policy
 - define
 - IPX 642
 - detail
 - IPX 641
 - modify
 - IPX 645
 - remove
 - IPX 648
 - summary
 - IPX 640
- script 124
- scripts for the Administration Console
 - examples 125
 - script command 124
- serial number 69, 129
- serial port (modem)
 - setting baud rate 95
- server information 116
- servers
 - defining static IPX 624
 - displaying static IPX 622
 - table for 622
- service levels
 - conforming packets 711
 - default 715
 - RSVP 728, 730
- services for event logging 752
- Simple Network Time Protocol (SNTP) 116 to 122
- size, burst 712
- SMT (Station Management)
 - lerAlarm value 242
 - lerCutoff value 243
- snapshot feature 86
- sniffer 755
- SNMP (Simple Network Management Protocol)
 - agent 189
 - community strings 192
 - display 189
 - displaying configurations 191
 - trap reporting
 - flushing addresses 197
- SNMP trap
 - addressThresholdEvent 257
- soft restarts 570
- software
 - backup NV data 107
 - build date and time 69, 129
 - version 69
- source address 480
 - advancedPing option 476
 - advancedTraceRoute option 480
 - traceRoute option 182
- source IP address for QoS classifiers 695, 702, 710, 718
- source IP address mask 695, 702
- source socket verification
 - AppleTalk 681
- speed, Ethernet ports 212, 213
- split horizon 457
- SRF (Status Report Frames)
 - and FDDI stations 224, 228
 - and lerAlarm 242
- state
 - STP mode (VLANs) 365
- state of IP interface 395
- static routes
 - defining for IP 422
 - defining for IPX 618
- static servers
 - defining IPX 624
- statistics
 - DDP (Datagram Delivery Protocol) 683
 - displaying IP, UDP, and ICMP 186
 - Ethernet 204, 208
 - general IP 482

ICMP (Internet Control Message Protocol) 187, 483
 IP interface 418
 IPX forwarding 653
 IPX interface 655
 IPX RIP 651
 IPX SAP 652
 NBP (Name Binding Protocol) 686
 OSPF (Open Shortest Path First) 603
 OSPF soft restart 570
 QoS (Quality of Service) interval for 738
 QoS receive 739
 QoS transmit 741
 RTMP (Routing Table Maintenance Protocol) 684
 trunk 305, 329
 UDP (User Datagram Protocol) 187, 483
 VLAN (virtual LAN) 341
 ZIP (Zone Information Protocol) 685
 statistics, AppleTalk protocol 683 to 686
 STP (Spanning Tree Protocol)
 stpMode 365
 stub default metrics, OSPF 571 to 573
 subnet masks
 defining 404, 406
 displaying 399
 for VLANs 349
 summary information
 trunk 301, 304, 328
 VLAN 338
 system baseline display 90
 system baseline set 91
 system console access 73
 system console webAccess 71
 system diagErrLog 115
 system ID 69
 system information
 displaying 69, 129
 system name
 displaying 69
 setting 101, 135, 136
 system reboot 123
 system serial port 93, 96
 system snmp define 117
 system snmp display 116
 system snmp modify 118
 system snmp pollInterval 121
 system snmp remove 119
 system snmp state 120
 system snmp timezone 104, 106, 121
 system snmp tolerance 122
 system up time 69

T

T_Opr 232
 tagging, VLAN
 defining 345, 352
 displaying 341
 modifying 355, 360
 specifying 347, 352
 tags, priority
 for controls 712
 for excess packets 736
 for nonflow classifiers 697, 703, 705
 TCMP (Trunk Control Message Protocol)
 displaying state 301, 304, 305, 328, 329
 technical support
 3Com Knowledgebase Web Services 767
 3Com URL 767
 bulletin board service 768
 fax service 769
 network suppliers 769
 product repair 771
 telnet
 rebooting the system 123
 terminal port
 port speed 93, 96
 terminal speed 93
 terminalSpeed command (system serialPort) 93, 96
 terminate a Telnet session 126
 TFTP (Trivial File Transfer Protocol) 87
 time
 displaying module 137
 time in service 69
 tolerance threshold 122
 total reservable bandwidth 725, 728
 traceRoute
 port number 182, 480
 source address 182
 using 478
 traceroute, IP multicast 528
 transmit statistics, QoS 741
 trap reporting
 adding 195
 flushing addresses 197
 modifying 195
 T-Req 232
 triggered updates
 RIP 631
 SAP 639
 trunk groups
 supported 314
 trunking
 and VLANs 346
 definition 299
 overview 299

trunks
 defining 310, 315
 definition 299
 detail information 305, 329
 maximum ports 316
 names 316
 removing 318
 resources 318
 sample definition 311
 summary information 304, 328
 trusted IP clients 77 to 82
 ttl (time to live)
 advancedTraceRoute 480
 example 185
 type of module 69

detail information and statistics 341
 displaying summary information 338
 errors 406
 interface index 401, 607
 modifying (Layer 2 devices) 360
 modifying (Layer 3 devices) 355, 358
 removing 363
 setting allOpen or allClosed mode 364
 setting ignore STP mode 365
 trunking 346
 VRRP (Virtual Router Redundancy Protocol)
 defining 492
 enabling or disabling 499
 introduction 485
 modifying 495
 removing 495, 498

U

UDP (User Datagram Protocol) Helper
 overlapped IP interfaces 447 to 449
 port and IP forwarding addresses 444
 UDP Helper
 BOOTP 442
 UDP port number
 advancedTraceRoute 480
 traceRoute 182
 UDP statistics 187, 483
 unspecified protocol type 348
 updates
 RIP triggered 631
 SAP triggered 639
 URL 767
 user configuration information 116

V

values 252
 default 41
 entering in command strings 41
 vi editor 124
 VID (VLAN ID) 341
 range 347, 352
 virtual links, OSPF 574 to 589
 VLAN interface 356
 VLAN interface index 356, 360
 specifying for ignore STP mode 365
 used to delete VLANs 363
 VLANs
 bridge VLAN commands
 modify 316
 VLANs (virtual LANs) 337
 defining for Layer 2 devices 352
 defining for Layer 3 devices 345

W

wait
 advancedPing option 475
 advancedTraceRoute option 480
 Web Management
 access 71, 72
 applications 30
 World Wide Web (WWW) 767

Z

ZIP (Zone Information Protocol) 685
 zones 677, 678

