

Configuring Switching Paths

This chapter describes how to configure switching paths on Cisco IOS devices. It provides configuration guidelines for switching paths and tuning guidelines. It also provides an overview of NetFlow. For documentation of switching commands used in this chapter, refer to the *Cisco IOS Switching Services Command Reference*. To find documentation of other commands, you can use the master indexes or search online.

Fast Switching Configuration Task List

Fast switching allows higher throughput by switching a packet using a cache created by the initial packet sent to a particular destination. Destination addresses are stored in the high-speed cache to expedite forwarding. Routers offer better packet-transfer performance when fast switching is enabled. Fast switching is enabled by default on all interfaces that support fast-switching.

Use the commands in these sections to configure appropriate fast-switching features.

- Enable AppleTalk Fast Switching
- Enable IP Fast Switching
- Enable Fast Switching on the Same IP Interface
- Enable Fast Switching of IPX Directed Broadcast Packets
- Disable Banyan VINES Fast Switching
- Enable Fast Switching of IPX Directed Broadcast Packets

Fast Switching is not supported for the X.25 encapsulations.

Enable AppleTalk Fast Switching

AppleTalk access lists are automatically fast switched. Access list fast switching improves the performance of AppleTalk traffic when access lists are defined on an interface. Refer to the “Configuring AppleTalk” chapter in the *Network Protocols Configuration Guide, Part 2* for guidelines on creating and using access lists and configuring AppleTalk.

Enable IP Fast Switching

Fast switching involves the use of a high-speed switching cache for IP routing. Destination IP addresses are stored in the high-speed cache to expedite packet forwarding. In some cases, fast switching is inappropriate, such as when slow-speed serial links (64K and below) are being fed from

higher-speed media such as T1 or Ethernet. In such a case, disabling fast switching can reduce the packet drop rate to some extent. Fast switching allows outgoing packets to be load balanced on a *per-destination* basis.

To enable or disable fast switching, use either of the following commands in interface configuration mode:


Command	Purpose
ip route-cache	Enable fast switching (use of a high-speed route cache for IP routing).
no ip route-cache	Disable fast switching and enable load balancing on a per-packet basis.

Enable Fast Switching on the Same IP Interface

You can enable IP fast switching when the input and output interfaces are the same interface. This normally is not recommended, though it is useful when you have partially meshed media such as Frame Relay. You could use this feature on other interfaces, although it is not recommended because it would interfere with redirection.

Figure 7 illustrates a scenario where this is desirable. Router A has a data link connection identifier (DLCI) to Router B, and Router B has a DLCI to Router C. There is no DLCI between Routers A and C; traffic between them must go in and out of Router B through the same interface.

Figure 7 IP Fast Switching on the Same Interface



To allow IP fast switching on the same interface, use the following command in interface configuration mode:

Command	Purpose
ip route-cache same-interface	Enable the fast switching of packets out of the same interface on which they arrived.

Enable Fast Switching of IPX Directed Broadcast Packets

By default, Cisco IOS software switches packets that have been directed to the broadcast address. To enable fast switching of these IPX-directed broadcast packets, use the following command in global configuration mode:

Command	Purpose
ipx broadcast-fastswitching	Enable fast switching of IPX directed broadcast packets.

Enable SMDS Fast Switching

SMDS fast switching of IP, IPX, and AppleTalk packets provides faster packet transfer on serial links with speeds above 56 kbps. Use fast switching if you use high-speed, packet-switched, datagram-based WAN technologies such as Frame Relay offered by service providers.

By default, SMDS fast switching is enabled.

To re-enable fast switching, if it has been disabled, use the following commands in interface configuration mode:

Step	Command	Purpose
1	interface <i>type number</i>	Define the type and unit number of the interface, and enter interface configuration mode.
2	encapsulation smds	Set SMDS encapsulation.
3	ip route-cache	Enable the interface for IP fast switching.
4	ipx route-cache	Enable the interface for IPX fast switching.
5	appletalk route-cache	Enable the interface for AppleTalk fast switching.

Disable Fast Switching for Troubleshooting

Fast switching uses a cache created by previous packets to achieve a higher packet throughput. Packet transfer performance is generally better when fast switching is enabled. Fast switching also provides load sharing on a per-packet basis.

By default, fast switching is enabled on all interfaces that support fast switching. However, you may want to disable fast switching to save memory space on interface cards and to help avoid congestion when high-bandwidth interfaces are writing large amounts of information to low-bandwidth interfaces. This is especially important when using rates slower than T1.

Fast switching is not supported on serial interfaces using encapsulations other than HDLC.

Note Turning off fast switching increases system overhead.

For some diagnostics such as debugging and packet-level tracing, you will need to disable fast switching. If fast switching is running, you will not see packets unless they pass through the route processor. Packets would otherwise be switched on the interface. You might want to turn off fast switching temporarily and bypass the route processor while you are trying to capture information.

This section includes these topics:

- Disable AppleTalk Fast Switching
- Disable Banyan VINES Fast Switching
- Disable DECnet Fast Switching
- Disable IPX Fast Switching
- Disable ISO CLNS Fast Switching Through the Cache
- Disable XNS Fast Switching

Disable AppleTalk Fast Switching

To disable AppleTalk fast-switching on an interface, use the following command in interface configuration mode:

Command	Purpose
no appletalk route-cache	Disable AppleTalk fast switching.

Disable Banyan VINES Fast Switching

Fast switching is enabled by default on all interfaces on which it is supported.

To disable fast switching on an interface, use the following command in interface configuration mode:

Command	Purpose
no vines route-cache	Disable fast switching.

Disable DECnet Fast Switching

By default, Cisco's DECnet routing software implements fast switching of DECnet packets.

To disable fast switching of DECnet packets, use the following command in interface configuration mode:

Command	Purpose
no decnet route-cache	Disable fast switching of DECnet packets on a per-interface basis.

Disable IPX Fast Switching

To disable IPX fast switching, use the following command in interface configuration mode:

Command	Purpose
no ipx route-cache	Disable IPX fast switching.

Disable ISO CLNS Fast Switching Through the Cache

ISO CLNS fast switching through the cache is enabled by default for all supported interfaces. To disable fast switching, use the following command in interface configuration mode:

Command	Purpose
no clns route-cache	Disable fast switching.

Note The cache still exists and is used after the **no clns route-cache** interface configuration command is used; the software just does not do fast switching through the cache.

Disable XNS Fast Switching

To disable XNS fast switching on an interface, use the following command in interface configuration mode:

Command	Purpose
<code>no xns route-cache</code>	Disable XNS fast switching.

Control the Route Cache

The high-speed route cache used by IP fast switching is invalidated when the IP routing table changes. By default, the invalidation of the cache is delayed slightly to avoid excessive CPU load while the routing table is changing. To control the route cache, use the appropriate commands in these sections:

- Control Route Cache Invalidation for IP
- Display System and Network Statistics
- Adjust the Route Cache for IPX
- Pad Odd-Length IPX Packets

Control Route Cache Invalidation for IP

To control route cache invalidation, use the following commands in global configuration mode as needed for your network:

Command	Purpose
<code>no ip cache-invalidate-delay</code>	Allow immediate invalidation of the cache.
<code>ip cache-invalidate-delay [minimum maximum quiet threshold]</code>	Delay invalidation of the cache.

Note This task normally should not be necessary. It should be performed only under the guidance of technical staff. Incorrect configuration can seriously degrade the performance of your router.

Display System and Network Statistics

You can display the contents of IP routing tables and caches. The resulting information can be used to determine resource utilization and to solve network problems.

Use the following command in privileged EXEC mode:

Command	Purpose
<code>show ip cache [prefix mask] [type number]</code>	Display the routing table cache used to fast switch IP traffic.

Adjust the Route Cache for IPX

Adjusting the route cache allows you to control the size of the route cache, reduce memory consumption, and improve router performance. You accomplish these tasks by controlling the route cache size and invalidation. The following sections describe these optional tasks:

- Control IPX Route Cache Size
- Control IPX Route Cache Invalidation

Control IPX Route Cache Size

You can limit the number of entries stored in the IPX route cache to free up router memory and aid router processing.

Storing too many entries in the route cache can use a significant amount of router memory, causing router processing to slow. This situation is most common on large networks that run network management applications for NetWare.

For example, if a network management station is responsible for managing all clients and servers in a very large (greater than 50,000 nodes) Novell network, the routers on the local segment can become inundated with route cache entries. You can set a maximum number of route cache entries on these routers to free up router memory and aid router processing.

To set a maximum limit on the number of entries in the IPX route cache, use the following command in global configuration mode:

Command	Purpose
ipx route-cache max-size <i>size</i>	Set a maximum limit on the number of entries in the IPX route cache.

If the route cache has more entries than the specified limit, the extra entries are not deleted. However, they may be removed if route cache invalidation is in use. See the “Control IPX Route Cache Invalidation” section in this chapter for more information on invalidating route cache entries.

Control IPX Route Cache Invalidation

You can configure the router to invalidate fast switch cache entries that are inactive. If these entries remain invalidated for one minute, the router purges the entries from the route cache.

Purging invalidated entries reduces the size of the route cache, reduces memory consumption, and improves router performance. Purging entries also helps ensure accurate route cache information.

You specify the period of time that valid fast switch cache entries must be inactive before the router invalidates them. You can also specify the number of cache entries that the router can invalidate per minute.

To configure the router to invalidate fast switch cache entries that are inactive, use the following command in global configuration mode:

Command	Purpose
ipx route-cache inactivity-timeout <i>period [rate]</i>	Invalidate fast switch cache entries that are inactive.

When you use the **ipx route-cache inactivity-timeout** command with the **ipx route-cache max-size** command, you can ensure a small route cache with fresh entries.

Pad Odd-Length IPX Packets

Some IPX end hosts accept only even-length Ethernet packets. If the length of a packet is odd, the packet must be padded with an extra byte so that end host can receive it. By default, Cisco IOS pads odd-length Ethernet packets.

However, there are cases in certain topologies where non-padded Ethernet packets are being forwarded onto a remote Ethernet network. Under specific conditions, you can enable padding on intermediate media as a temporary workaround for this problem. Note that you should perform this task only under the guidance of a customer engineer or other service representative.

To enable the padding of odd-length packets, use the following commands in interface configuration mode:

Step	Command	Purpose
1	no ipx route-cache	Disable fast switching.
2	ipx pad-process-switched-packets	Enable the padding of odd-length packets.

