

Configuring Switching Paths

This chapter describes switching paths that can be configured on Cisco IOS devices. It provides an overview of switching methods, 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*. For documentation of other commands you can use the master indexes or search online.


Overview of Basic Router Platform Architecture and Processes

To understand how switching works, it helps to first understand the basic router architecture and where various processes occur in the router.

Fast switching is enabled by default on all interfaces that support fast switching. If you have a situation where you need to disable fast switching and fall back to the process-switching path, understanding how various processes affect the router and where they occur will help you determine your alternatives. This is especially true when you are troubleshooting traffic problems or need to process packets that require special handling. Some diagnostic or control resources are not compatible with fast switching or come at the expense of processing and switching efficiency. Understanding the effects of those resources can help you minimize their effect on network performance.

Figure 2 illustrates a possible internal configuration of a Cisco 7500 series router. In this configuration, the Cisco 7500 series router has an integrated Route/Switch Processor (RSP) and uses *route caching* to forward packets. The Cisco 7500 series router also uses Versatile Interface Processors (VIPs), a RISC-based interface processor that receives and caches routing information from the RSP. Using the routing cache, the VIP card makes switching decisions locally, relieving the RSP of involvement and speeding overall throughput. This type of switching is called *distributed switching*. Multiple VIP cards can be installed in one router.

Figure 2 Basic Router Architecture



Cisco Routing and Switching Processes

The routing, or forwarding, function comprises two interrelated processes to move information in the network.

- Making a routing decision by routing
- Moving packets to the next-hop destination by switching

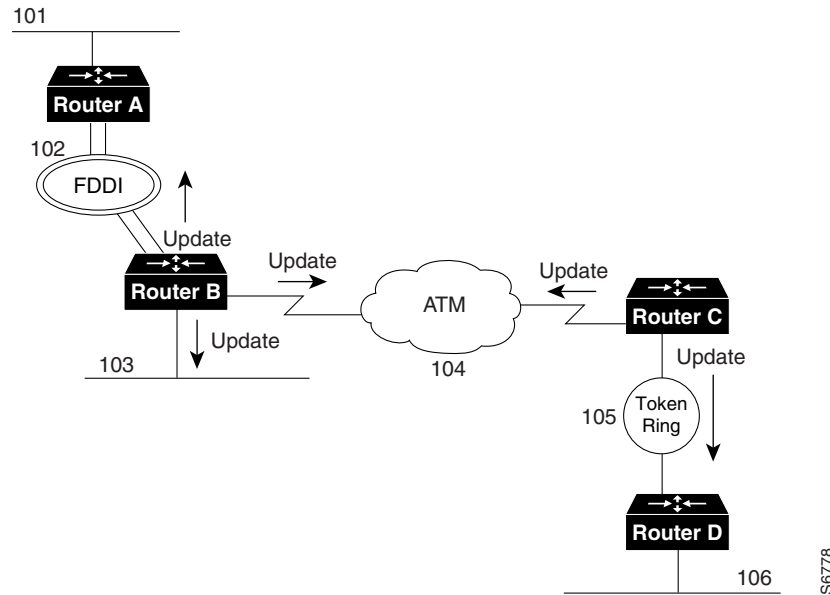
Cisco IOS platforms perform both routing and switching, and there are several types of each.

Routing

The routing process assesses the source and destination of traffic based on knowledge of network conditions. Routing functions identify the best path to use for moving the traffic to the destination out one or more of the router interfaces. The routing decision is based upon a variation of criteria such as link speed, topological distance, and protocol. Each separate protocol maintains its own routing information.

Routing is more processing intensive and has higher latency than switching as it determines path and next-hop considerations. The first packet routed requires a lookup in the routing table to determine the route. The route cache is populated after the first packet is routed by the route-table lookup. Subsequent traffic for the same destination is switched using the routing information stored in the route cache. Figure 3 illustrates the basic routing process.

Figure 3 The Routing Process




A router sends routing updates out each of its interfaces that are configured for a particular protocol. It also receives routing updates from other attached routers. From these received updates and its knowledge of attached networks, it builds a map of the network topology.

Switching


Through the switching process, the router determines the next hop toward the destination address. Switching moves traffic from an input interface to one or more output interfaces. Switching is optimized and has lower latency than routing because it can move packets, frames, or cells from buffer to buffer with simpler determination of the source and destination of the traffic. It saves resources because it does not involve extra lookups. Figure 4 illustrates the basic switching process.

Figure 4 The Switching Process



In Figure 4, packets are received on the Fast Ethernet interface and destined for the FDDI interface. Based on information in the packet header and destination information stored in the routing table, the router determines the destination interface. It looks in the protocol's routing table to discover the destination interface that services the destination address of the packet.

The destination address is stored in tables such as ARP tables for IP and AARP table for AppleTalk. If there is no entry for the destination, the router will either drop the packet (and inform the user if the protocol provides that feature), or it must discover the destination address by some other address resolution process, such as through the ARP protocol. Layer 3 IP addressing information is mapped to the Layer 2 MAC address for the next hop. Figure 5 illustrates the mapping that occurs to determine the next hop.



Basic Switching Paths

Basic switching paths are

- Process Switching
- Fast Switching
- Optimum Switching
- Distributed Switching
- NetFlow

Process Switching

In process switching the first packet is copied to the system buffer. The router look up the Layer 3 network address in the routing table and initializes the fast-switch cache. The frame is rewritten with the destination address and sent to the exit interface that services that destination. Subsequent packets for that destination are sent by the same switching path. The route processor computes the cyclical redundancy check (CRC).

Fast Switching

When packets are fast switched, the first packet is copied to packet memory and the destination network or host is found in the fast-switching cache. The frame is rewritten and sent to the exit interface that services the destination. Subsequent packets for the same destination use the same switching path. The interface processor computes the CRC.

Optimum Switching

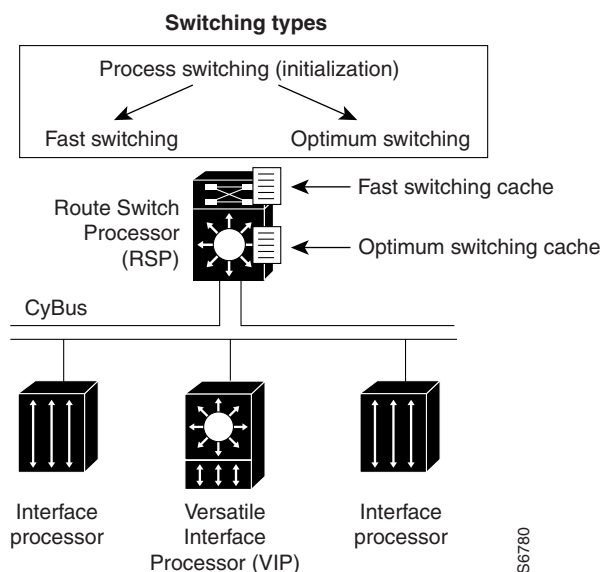
Optimum switching is similar to fast switching, but is faster. The first packet is copied to packet memory and the destination network or host is found in the optimum-switching cache. The frame is rewritten and sent to the exit interface that services the destination. Subsequent packets for the same destination use the same switching path. The interface processor computes the CRC.

Optimum switching is enabled by default on Cisco 7500 series routers; it must be disabled for debugging.

Distributed Switching

Switching becomes more efficient the closer to the interface the function occurs. In distributed switching, the switching process occurs on VIP and other interface cards that support switching. For model numbers and hardware compatibility information, refer to the *Cisco Product Catalog*. Figure 6 illustrates the distributed switching process on the Cisco 7500 series.

Figure 6 Distributed Switching on Cisco 7500 Series Routers



The VIP card installed in this router maintains a copy of the routing cache information needed to forward packets. Because the VIP card has the routing information it needs, it performs the switching locally, making the packet forwarding much faster. Router throughput is increased linearly based on the number of VIP cards installed in the router.

NetFlow

While in more recent IOS releases it is not a switching path, NetFlow enables you to collect the data required for flexible and detailed accounting, billing, and chargeback for network and application resource utilization. Accounting data can be collected for both dedicated line and dial-access accounting. NetFlow is supported over switched LAN or ATM backbones, allowing scalable inter-VLAN forwarding. NetFlow can be deployed at any location in the network. NetFlow is available on the Cisco 7200 Series and the Cisco 7500 Series.

NetFlow is configurable on a per-interface basis. You configure NetFlow using the *protocol route-cache flow* command.

NetFlow is described in "Configuring NetFlow" later in this publication.

Platform and Switching Path Correlation

Depending on the routing platform you are using, availability and default implementations of switching paths varies. Table 1 shows the correlation between Cisco IOS switching paths and routing platforms.

Table 1 Switching Paths on RSP-Based Routers

Switching Path	Cisco 7200	Cisco 7500	Comments	Configuration Command
Process switching	Yes	Yes	Initializes switching caches	no protocol route-cache
Fast switching	Yes	Yes	Default (except for IP)	<i>protocol route-cache</i>
Optimum switching	Yes	Yes	Default for IP	<i>protocol route-cache optimum</i>
Distributed switching	No	Yes	Using Second-Generation VIP line cards	<i>protocol route-cache distributed</i>

Understanding Features that Affect Performance

Performance is derived from the switching mechanism you are using. Some Cisco IOS features require special handling and cannot be switched until the additional processing they require has been performed. This special handling is not processing that the interface processors can do. Because these features require additional processing, they affect switching performance. These features include

- Queuing
- Random Early Detection
- Compression
- Filtering (using access lists)
- Encryption
- Accounting

Queuing

Queuing occurs when network congestion occurs. When traffic is moving well within the network, packets are sent as they arrive at the interface. Cisco IOS software implements four different queuing algorithms:

- First In, First Out (FIFO) Queuing – Packets are forwarded in the same order in which they arrive at the interface.
- Priority Queuing – Packets are forwarded based on an assigned priority. You can create priority lists and groups to define rules for assigning packets to priority queues.
- Custom Queuing – You can control a percentage of interface bandwidth for specified traffic by creating protocol queue lists and custom queue lists.
- Weighted Fair Queuing – Weighted Fair Queuing provides automatic traffic priority management. Low-bandwidth sessions have priority over high-bandwidth sessions and high-bandwidth sessions are assigned weights. Weighted Fair Queuing is the default for interfaces slower than 2.048 Mbps.

Random Early Detection

Random Early Detection is designed for congestion avoidance. Traffic is prioritized based on type of service (TOS), or precedence. This feature is available on T3, OC-3, and ATM interfaces.

Compression

Depending on the protocol you are using, various compression options are available in Cisco IOS software. Refer to the Cisco IOS configuration guide for the protocol you are using to see what compression options you have.

Filtering

You can define access lists to control access to or from a router for a number of services. You could, for example, define an access lists to prevent packets with a certain IP address from leaving a particular interface on a router. How access lists are used depends on the protocol. For information on access lists, refer to the appropriate Cisco IOS configuration guide for the protocol you are using.

Encryption

Encryption algorithms are applied to data to alter its appearance making it incomprehensible to those who are not authorized to see the data. For information about encryption features available with the Cisco IOS software, refer to the *Security Configuration Guide*.

Accounting

You can configure accounting features to collect network data related to resource usage. The information you collect (in the form of statistics) can be used for billing, chargeback, and planning resource usage. Refer to appropriate Cisco IOS configuration guide for the protocol you are using for information regarding accounting features you can use.

Configuring Fast Switching

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.

Perform the tasks 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, perform either of the following tasks in interface configuration mode:


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

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, perform this task in interface configuration mode:

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

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, perform the following task in global configuration mode:

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

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, perform the following tasks in interface configuration mode:

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

Disabling Fast Switching for Troubleshooting

Fast switching allows higher throughput by switching packets using a cache created by previous packets. 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, perform the following task in interface configuration mode:

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

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, perform the following task in interface configuration mode:

Task	Command
Disable fast switching.	no vines route-cache

Disable DECnet Fast Switching

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

To disable fast switching of DECnet packets, perform the following task in interface configuration mode:

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

Disable IPX Fast Switching

To disable IPX fast switching, perform the following task in interface configuration mode:

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

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, perform the following task in interface configuration mode:

Task	Command
Disable fast switching.	no clns route-cache

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, perform the following task in interface configuration mode:

Task	Command
Disable XNS fast switching.	no xns route-cache

Disabling Optimum Switching for Troubleshooting

Optimum switching is enabled by default for IP on Ethernet, FDDI, serial interfaces on the Cisco 7500 series, and on all ATM port adapter interfaces. On serial interfaces, it is supported for HDLC encapsulation only. If you want to use fast-switching or process switching, disable optimum switching by performing the following task in interface configuration mode:

Task	Command
Disable optimum switching.	no ip route-cache optimum

Optimum switching must be disabled for troubleshooting.

Controlling 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, perform the appropriate tasks 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, perform the following tasks in global configuration mode as needed for your network:

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

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.

Perform the following task in privileged EXEC mode:

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

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, complete this task in global configuration mode:

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

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, complete this task in global configuration mode:

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

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, perform the following tasks in interface configuration mode:

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

