



Configuring Cisco Express Forwarding

This chapter describes the required and optional tasks for configuring Cisco Express Forwarding (CEF) and distributed CEF (dCEF). For a complete description of CEF commands used in this chapter, refer to the *Cisco IOS Switching Services Command Reference*. For documentation on other commands that appear in this chapter, you can use the command reference master index or search online.

Cisco Express Forwarding Configuration Task List

To configure CEF, complete the tasks in the following sections:

- Enabling CEF or dCEF (Required)
- Configuring Load Balancing for CEF (Optional)
- Configuring Network Accounting for CEF (Optional)
- Configuring Distributed Tunnel Switching for CEF (Optional)

Enabling CEF or dCEF

Enable CEF when your router has interface processors that do not support dCEF.

To enable CEF, use the following command in global configuration mode:

Command	Purpose
<code>ip cef</code>	Enables standard CEF operation.

Enable dCEF when you want your line cards to perform express forwarding so that the route processor can handle routing protocols or switch packets from legacy interface processors.



Note

On the Cisco 12000 series routers, dCEF is enabled by default. The command to enable dCEF is not available. Also, the configuration file does not indicate that dCEF is enabled on the router.

To enable or disable dCEF operation, use one of the following commands in global configuration mode:

Command	Purpose
<code>ip cef distributed</code>	Enables dCEF operation.
<code>no ip cef distributed</code>	Disables dCEF operation.

When you enable CEF or dCEF globally, all interfaces that support CEF are enabled by default. If you want to turn off CEF or dCEF on a particular interface, you can do so.

To disable CEF or dCEF on an interface, use the following command in interface configuration mode:

Command	Purpose
<code>no ip route-cache cef</code>	Disables CEF operation on the interface.

When you disable CEF or dCEF, Cisco IOS software switches packets using the next fastest switching path. In the case of dCEF, the next fastest switching path is CEF on the route processor.

If you have disabled CEF or dCEF operation on an interface and want to reenoble it, you can do so by using the **ip route-cache cef** command from interface configuration mode.



Note

On the Cisco 12000 series routers, you must not disable dCEF on an interface.

Configuring Load Balancing for CEF

CEF load balancing is based on a combination of source and destination packet information; it allows you to optimize resources by distributing traffic over multiple paths for transferring data to a destination. You can configure load balancing on a per-destination or per-packet basis. Load-balancing decisions are made on the outbound interface. When you configure load balancing, configure it on outbound interfaces.

You can configure two types of load balancing for CEF by performing the following optional tasks:

- Configuring Per-Destination Load Balancing
- Configuring Per-Packet Load Balancing

Configuring Per-Destination Load Balancing

Per-destination load balancing is enabled by default when you enable CEF. To use per-destination load balancing, you do not perform any additional tasks once you enable CEF.

Per-destination load balancing allows the router to use multiple paths to achieve load sharing. Packets for a given source-destination host pair are guaranteed to take the same path, even if multiple paths are available. Traffic destined for different pairs tend to take different paths. Per-destination load balancing is enabled by default when you enable CEF, and is the load balancing method of choice for most situations.

Because per-destination load balancing depends on the statistical distribution of traffic, load sharing becomes more effective as the number of source-destination pairs increase.

You can use per-destination load balancing to ensure that packets for a given host pair arrive in order. All packets for a certain host pair are routed over the same link (or links).

Disabling Per-Destination Load Balancing

Typically, you would disable per-destination load balancing when you want to enable per-packet load balancing.

To disable per-destination load balancing, use the following command in interface configuration mode:

Command	Purpose
<code>no ip load-sharing per-destination</code>	Disables per-destination load balancing.

Configuring Per-Packet Load Balancing

Per-packet load balancing allows the router to send successive data packets over paths without regard to individual hosts or user sessions. It uses the round-robin method to determine which path each packet takes to the destination. Per-packet load balancing ensures balancing over multiple links.



Note

Per-packet load balancing via CEF is not supported on Engine 2 Gigabit Switch Router (GSR) line cards (LCs).

Path utilization with per-packet load balancing is good, but packets for a given source-destination host pair might take different paths. Per-packet load balancing could introduce reordering of packets. This type of load balancing would be inappropriate for certain types of data traffic (such as voice traffic over IP) that depend on packets arriving at the destination in sequence.

Use per-packet load balancing to help ensure that a path for a single source-destination pair does not get overloaded. If the bulk of the data passing through parallel links is for a single pair, per-destination load balancing will overload a single link while other links have very little traffic. Enabling per-packet load balancing allows you to use alternate paths to the same busy destination.

To enable per-packet load balancing, use the following command in interface configuration mode:

Command	Purpose
<code>ip load-sharing per-packet</code>	Enables per-packet load balancing.



Note

If you want to enable per-packet load sharing to a particular destination, all interfaces that can forward traffic to the destination must be enabled for per-packet load sharing.

Configuring Network Accounting for CEF

You might want to collect statistics to better understand CEF patterns in your network. For example, you might want to collect information such as the number of packets and bytes switched to a destination or the number of packets switched through a destination.

You can configure network accounting for CEF by performing the following optional tasks:

- Enabling Network Accounting for CEF
- Viewing Network Accounting Information

Enabling Network Accounting for CEF

To collect network accounting information for CEF, use one of the following commands in global configuration mode:

Command	Purpose
<code>ip cef accounting per-prefix</code>	Enables the collection of the number of packets and bytes express forwarded to a destination (or prefix).
<code>ip cef accounting non-recursive</code>	Enables the collection of the number of packets express forwarded through a destination.

When you enable network accounting for CEF from global configuration mode, accounting information is collected at the route processor when CEF mode is enabled. When network accounting is enabled for dCEF, information is collected at the line cards.

Viewing Network Accounting Information

You can view the collected accounting information. To do so, use the following command in EXEC mode:

Command	Purpose
<code>show ip cef</code>	Displays the collected accounting information.

Configuring Distributed Tunnel Switching for CEF

CEF supports distributed tunnel switching, such as GRE tunnels. Distributed tunnel switching is enabled automatically when you enable CEF or dCEF. You do not perform any additional tasks to enable distributed tunnel switching once you enable CEF or dCEF.