Cisco Express Forwarding

Feature Summary

Cisco Express Forwarding (CEF) is advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns, such as the Internet, on networks characterized by intensive Web-based applications, or interactive sessions.

Although you can use CEF in any part of a network, it is designed for high-performance, highly resilient Layer 3 IP backbone switching.

Benefits

CEF offers these benefits:

- Improved performance—CEF is less CPU-intensive than fast or optimum switching route caching. More CPU processing power can be dedicated to Layer 3 services such as quality of service (QoS) and encryption.
- Scalability—CEF offers full switching capacity at each line card when dCEF mode is active.
- Resilience—CEF offers unprecedented level of switching consistency and stability in large dynamic networks. In dynamic networks, fast switching cache entries are frequently invalidated due to routing changes. These changes can cause traffic to be process switched using the routing table, rather than fast switched using the route cache. Because the FIB lookup table contains all known routes that exist in the routing table, it eliminates route cache maintenance and the fast switch/process switch forwarding scenario. CEF can switch traffic more efficiently than typical demand caching schemes.

List of Terms

Cisco Express Forwarding (CEF) Layer 3 switching technology. CEF can also refer to central CEF mode, one of the two modes of CEF operation that enables a route processor to perform express forwarding.

Distributed CEF (dCEF) One of two modes of CEF operation that enables line cards to perform the express forwarding between port adapters.

Forwarding Information Base (FIB) A component of CEF. It is the lookup table the router uses to make destination-based switching decisions during CEF operation. It maintains a mirror image of the forwarding information stored in the IP routing table.

Gigabit Switch Router (GSR) Also called the Cisco 12000 series router.

line card—A general term for an interface processor that can be used in a line of Cisco products. For example, a VIP is a line card for the Cisco 7500 series router.

prefix—Indicates the network address portion of an IP address is the network versus the host. It is a combination of network and mask. The mask indicates which bits are the network bits. The prefix is in the format *network/mask*. For example, 10.0.0.0/16 means that the first 16 bits of the IP address are masked, making them the network bits and the remaining bits the host bits. In this example, the network number is 10.0.

route processor—A general term for a processor module on either the Cisco 7000 or Cisco 7500 router.

Route Processor (RP) Processor module on the Cisco 7000 series routers that contains the CPU, system software, and most of the memory components that are used in the router.

Route Switch Processor (RSP) Processor module used in the Cisco 7500 series routers that integrates the functions of the RP and the Switch Processor (SP). The SP acts as the administrator for all data bus activities on a Cisco 7000 series router.

Versatile Interface Processor (VIP) Interface card used in Cisco 7000 and Cisco 7500 series routers. The VIP provides multilayer switching and runs the Cisco IOS software. The most recent version of VIP is the VIP2.

Inter Process Communication (IPC) The mechanism that ensures synchronization of FIBs and adjacency tables on the RSP and line cards during dCEF mode.

Restrictions

Consider these items when implementing CEF in your network:

- The recommended minimum memory requirement in platforms carrying the current full Internet routing information is
  - 128 MB for the centralized route processor
  - 64 MB for each line card
- CEF cannot be run simultaneously with VIP-distributed fast switching in the same line card.
- By default, Cisco IOS software switches packets using the next fastest switching path (such as optimum, fast, or process switching) when CEF does not support a feature or encapsulation.
- CEF currently does not support these features:
  - Policy routing
  - Network Address Translation (NAT)
  - Access lists on the GSR
  - Multipoint PPP
  - SMDS
  - Token Ring
  - ATM dixie
  - Inter-Switch Link (ISL)

Platforms

This feature is supported on these platforms:

- Cisco 7000 series routers equipped with RSP7000
- Cisco 7200 series
- Cisco 7500 series
- Cisco 12000 series

CEF is supported on second-generation Versatile Interface Processor models VIP2-20, VIP2-40, and VIP2-50.

Depending on the router platform you are using for CEF, the exact hardware performing the switching might differ. Exactly where the switching occurs in the router depends on the router model and the hardware installed in the router. For example, on the Cisco 12000 series, all switching occurs on the line cards. On the Cisco 7500 series, packets can be switched concurrently by the RSP and VIP line cards. Each interface card runs its own express forwarding engine and maintains its own exact copy of the FIB table. Each card, independently, is able to forward packets, reducing the number of packets centrally switched by the route processor.

For the purpose of simplicity and consistency, discussions throughout this document use route processor to indicate the central processor and line card to indicate the interface processor or other line card.

For information about specific hardware requirements and compatibility, refer to the *Cisco Product Catalog* or the installation and configuration guide for the type of router you are using for CEF switching.

Supported MIBs and RFCs

None

Functional Description

Cisco Express Forwarding (CEF) is advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with large and dynamic traffic patterns, such as the Internet, on networks characterized by intensive Web-based applications, or interactive sessions.

Although you can use CEF in any part of a network, it is designed for high-performance, highly resilient Layer 3 IP backbone switching. For example, Figure 1 shows CEF being run on Cisco 12000 series Gigabit Switch Routers (GSRs) at aggregation points at the core of a network where traffic levels are dense and performance is critical.

Figure 1 Cisco Express Forwarding

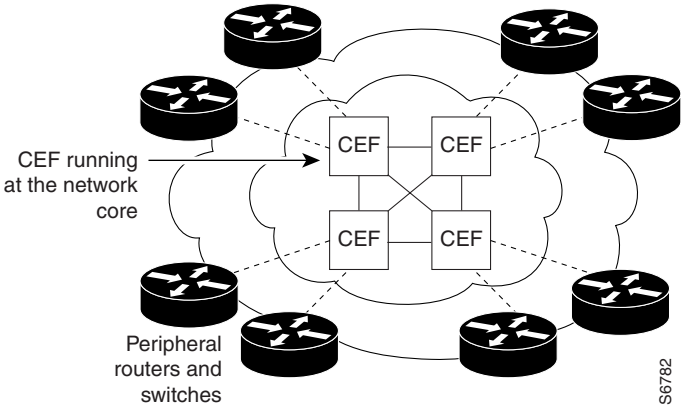

In a typical high-capacity internet service provider environment, Cisco 12012 GSRs as aggregation devices at the core of the network support links to Cisco 7500 series routers or other feeder devices. CEF in these platforms at the network core provides the performance and scalability needed to respond to continued growth and steadily increasing network traffic. CEF is a distributed switching mechanism that scales linearly with the number of interface cards and bandwidth installed in the router.

CEF Components

Information conventionally stored in a route cache is stored in several data structures for CEF switching. The data structures provide optimized lookup for efficient packet forwarding. The two main components of CEF operation are the

- Forwarding Information Base
- Adjacency Tables

Forwarding Information Base

CEF uses a Forwarding Information Base (FIB) to make IP destination prefix-based switching decisions. The FIB is conceptually similar to a routing table or information base. It maintains a mirror image of the forwarding information contained in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB maintains next-hop address information based on the information in the IP routing table.

Because there is a one-to-one correlation between FIB entries and routing table entries, the FIB contains all known routes and eliminates the need for route cache maintenance that is associated with earlier switching paths such as fast switching and optimum switching.

Adjacency Tables

Network nodes in the network are said to be adjacent if they can reach each other with a single hop across a link layer. In addition to the FIB, CEF uses adjacency tables to prepend Layer 2 addressing information. The adjacency table maintains Layer 2 next-hop addresses for all FIB entries.

Adjacency Discovery

The adjacency table is populated as adjacencies are discovered. Each time an adjacency entry is created (such as through the ARP protocol), a link-layer header for that adjacent node is precomputed and stored in the adjacency table. Once a route is determined, it points to a next hop and corresponding adjacency entry. It is subsequently used for encapsulation during CEF switching of packets.

Adjacency Resolution

A route might have several paths to a destination prefix, such as when a router is configured for simultaneous load balancing and redundancy. For each resolved path, a pointer is added for the adjacency corresponding to the next-hop interface for that path. This mechanism is used for load balancing across several paths.

Adjacency Types That Require Special Handling

In addition to adjacencies associated with next-hop interfaces (host-route adjacencies), other types of adjacencies are used to expedite switching when certain exception conditions exist. When the prefix is defined, prefixes requiring exception processing are cached with one of the special adjacencies listed in Table 1.

Table 1 Adjacency Types for Exception Processing

This adjacency type...	Receives this processing...
Null adjacency	Packets destined for a Null0 interface are dropped. This can be used as an effective form of access filtering.
Glean adjacency	When a router is connected directly to several hosts, the FIB table on the router maintains a prefix for the subnet rather than for the individual host prefixes. The subnet prefix point to a glean adjacency. When packets need to be forwarded to a specific host, the adjacency database is gleaned for the specific prefix.
Punt adjacency	Features that require special handling or features that are not yet supported in conjunction with CEF switching paths are forwarded to the next switching layer for handling. Features that are not supported are forwarded to the next higher switching level.
Discard adjacency	Packets are discarded. This type of adjacency occurs only on the Cisco 12000 series routers.
Drop adjacency	Packets are dropped, but the prefix is checked.

Unresolved Adjacency

When a link-layer header is prepended to packets, FIB requires the prepend to point to an adjacency corresponding to the next hop. If an adjacency was created by FIB and not discovered through a mechanism, such as ARP, the Layer 2 addressing information is not known and the adjacency is considered incomplete. Once the Layer 2 information is known, the packet is forwarded to the route processor, and the adjacency is determined through ARP.

Supported Media

CEF currently supports ATM/AAL5snap, ATM/AAL5mux, ATM/AAL5nlpid, Frame Relay, Ethernet, FDDI, PPP, HDLC, and tunnels.

CEF Operation Modes

CEF can be enabled in one of two modes:

- Central CEF Mode
- Distributed CEF Mode

Central CEF Mode

When CEF mode is enabled, the CEF FIB and adjacency tables reside on the route processor, and the route processor performs the express forwarding. You can use CEF mode when line cards are not available for CEF switching or when you need to use features not compatible with distributed CEF switching.

Figure 2 shows the relationship between the routing table, FIB, and adjacency table during CEF mode.

Figure 2 CEF Mode

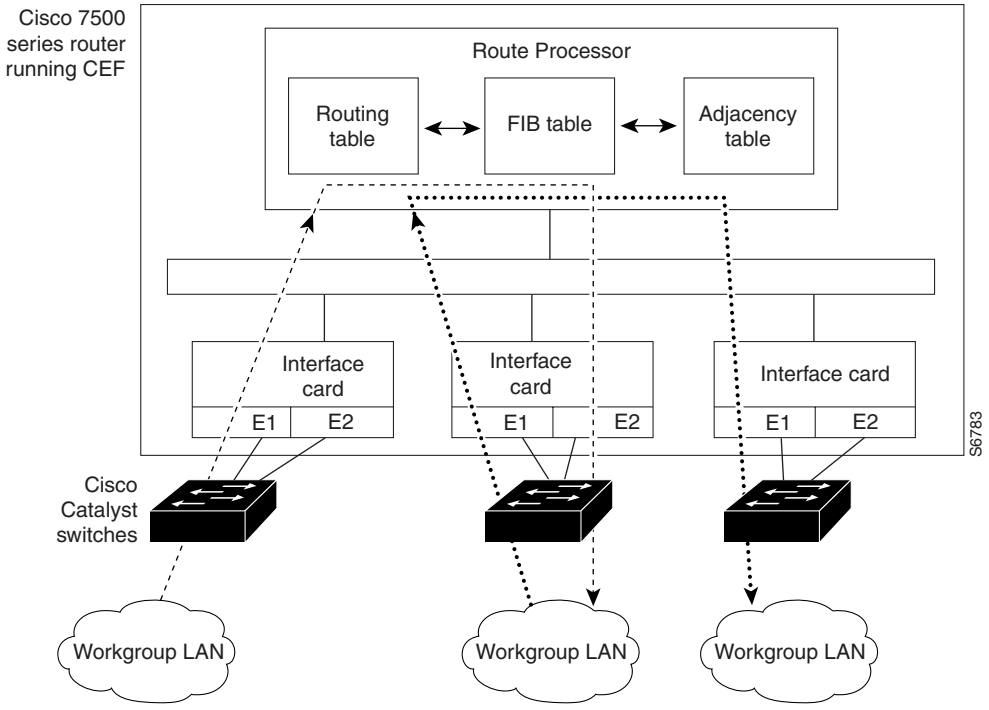

Figure 2 shows Cisco Catalyst switches forwarding traffic from workgroup LANs to a Cisco 7500 series router on the enterprise backbone running Cisco Express Forwarding. The route processor performs the express forwarding.

Distributed CEF Mode

When distributed CEF (dCEF) is enabled, line cards, such as VIP line cards or Gigabit Switch Router (GSR) line cards, maintain an identical copy of the FIB and adjacency tables. The line cards perform the express forwarding between port adapters, relieving the RSP of involvement in the switching operation.

dCEF uses an Inter Process Communication (IPC) mechanism to ensure synchronization of FIBs and adjacency tables on the route processor and line cards.

Figure 3 shows the relationship between the route processor and line cards when dCEF mode is active.

Figure 3 dCEF Mode

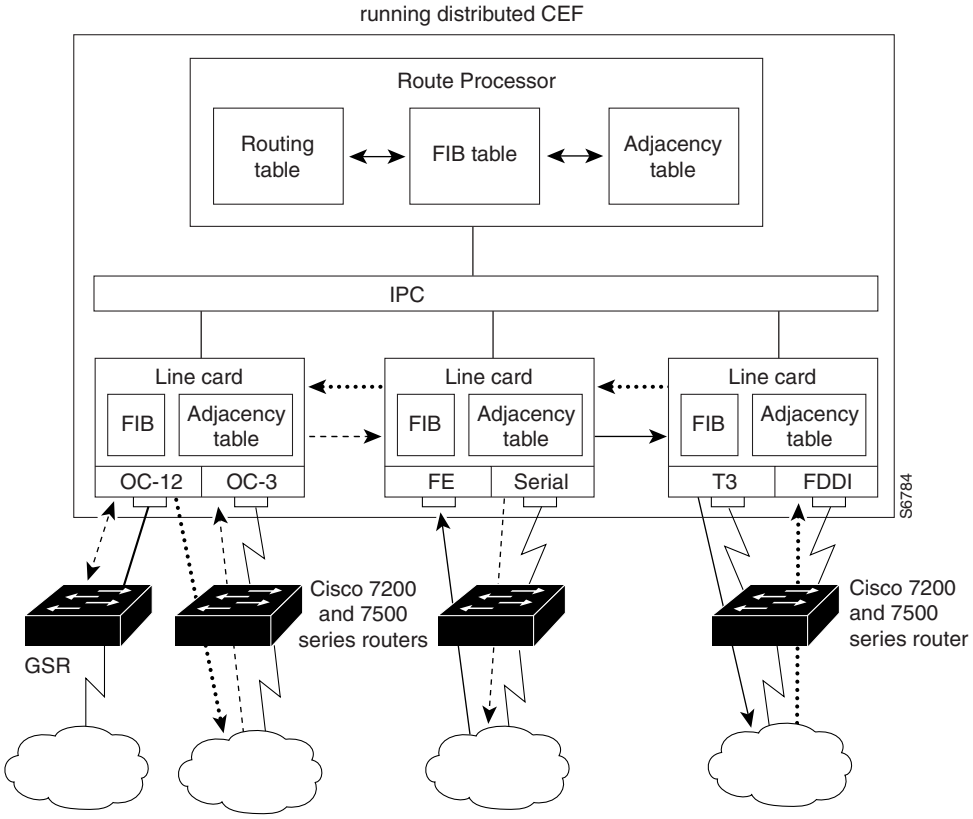

In this Cisco 12000 series router the line cards perform the switching. In other routers where you can mix various types of cards in the same router, it is possible that not all of the cards you are using support CEF. When a line card that does not support CEF receives a packet, the line card forwards the packet to the next higher switching layer (the route processor) or forwards the packet to the next hop for processing. This structure allows legacy interface processors to exist in the router with newer interface processors.

Notes

1 The Cisco 12000 series Gigabit Switch Routers operate only distributed CEF (dCEF mode).

2 Distributed CEF switching cannot be configured on the same VIP card as distributed fast switching.

3 Distributed CEF is not supported on Cisco 7200 series routers.

Additional Capabilities

In addition to configuring CEF and dCEF, you can also configure these features:

- Distributed CEF switching using access lists
- Distributed CEF switching of Frame Relay packets
- Distributed CEF switching during packet fragmentation
- Load balancing on a per destination or per packet basis
- Network accounting to gather byte and packet statistics
- Distributed CEF switching across IP tunnels

Configuration Tasks

This section describes the required and optional tasks for configuring CEF. The first task is required; all other tasks are optional.

- Enable and Disable CEF or dCEF
- Configure Load Balancing for CEF
- Configure Network Accounting for CEF
- Configure Distributed Tunnel Switching for CEF

Enable and Disable CEF or dCEF

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

To enable or disable CEF, perform one of the following tasks in global configuration mode:

Task	Command
Enable standard CEF operation.	**ip cef switch**
Disable standard CEF operation	**no ip cef switch**

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 list that dCEF is enabled on the router.

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

Task	Command
Enable dCEF operation.	**ip cef distributed switch**
Disable dCEF operation.	**no ip cef distributed switch**

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.

You might want to disable CEF or dCEF on a particular interface because that interface is configured with a feature that CEF or dCEF does not support. For example, policy routing and CEF cannot be used together. You might want one interface to support policy routing while the other interfaces support CEF. In this case, you would enable CEF globally, but disable CEF on the interface configured for policy routing, enabling all but one interface to express forward.

To disable CEF or dCEF on an interface, perform the following task in interface configuration mode:

Task	Command
Disable CEF operation on the interface.	**no ip route-cache cef**

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 reenable 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.

Configure 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.

Configure Load Balancing Task List

These sections describe how to configure each type of load balancing:

- Configure Per-Destination Load Balancing
- Configure Per-Packet Load Balancing

Configure Per-Destination Load Balancing

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).

Enable 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.

Disable 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, perform the following task in interface configuration mode:

Task	Command
Disable per-destination load balancing.	**no ip load-sharing per-destination**

Configure 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.

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, perform the following task in interface configuration mode:

Task	Command
Enable per-packet load balancing.	**ip load-sharing per-packet**

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

Configure 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 these optional tasks:

- Enable Network Accounting for CEF
- View Network Accounting Information

Enable Network Accounting for CEF

To collect network accounting information for CEF, perform the following tasks in global configuration mode:

Task	Command
Enable the collection of the number of packets and bytes express forwarded to a destination (or prefix).	**ip cef accounting per-prefix**
Enable the collection of the number of packets express forwarded through a destination.	**ip cef accounting non-recursive**

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 distributed CEF (dCEF), information is collected at the line cards.

View Network Accounting Information

You can then view the collected accounting information. To do so, perform the following task in EXEC mode:

Task	Command
Display the collected accounting information.	**show ip cef**

Configure 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.

Command Reference

This section documents new or modified commands. All other commands used with this feature are documented in the Cisco IOS Release 11.2 command references.

- **clear adjacency**
- **clear cef linecard**
- **clear ip cef prefix-statistics**
- **ip cef**
- **ip cef accounting**
- **ip load-sharing**
- **ip route-cache cef**
- **show adjacency**
- **show cef**
- **show cef interface**
- **show cef linecard**
- **show ip cef**

clear adjacency

To clear the Cisco Express Forwarding (CEF) adjacency table, use the **clear adjacency** EXEC command.

clear adjacency

Syntax Description

This command has no arguments or keywords.

Command Mode

EXEC

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

When you issue this command, entries in the adjacency table that resides on the route processor are removed and then repopulated. During repopulation, Layer 2 next hop information is reevaluated.

With dCEF mode, the adjacency tables that reside on line cards are always synchronized to the adjacency table that resides on the route processor. Therefore, clearing the adjacency table on the route processor using the **clear adjacency** command also clears the adjacency tables on the line cards; all changes are propagated to the line cards.

Related Commands

show adjacency

clear cef linecard

To clear Cisco Express Forwarding (CEF) information from line cards, use the **clear cef linecard** EXEC command.

clear cef linecard [*slot-number*] [**adjacency** | **interface** | **prefix**]

Syntax Description

slot-number	(Optional) Line card slot number to clear. When you omit this argument, all line card slots are cleared.
adjacency	(Optional) Clears line card adjacency tables and rebuilds adjacency for the specified line card.
interface	(Optional) Clears line card interface information and re-creates the interface information for the specified line card.
prefix	(Optional) Clears line card prefix tables and starts rebuilding the FIB table.

Command Mode

EXEC

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

This command is available only on routers with line cards. This command clears CEF information only on the line cards; CEF information on the route processor is not affected.

Once you clear CEF information from line cards, the corresponding information from the route processor is propagated to the line cards. The Inter Process Communication (IPC) ensures that CEF information on the route processor matches the CEF information on the line cards.

Related Commands

show cef linecard

clear ip cef prefix-statistics

To clear Cisco Express Forwarding (CEF) counters by resetting the packet and byte count to zero (0), use the **clear ip cef prefix-statistics** EXEC command.

clear ip cef {*network* [*mask*] | *****} **prefix-statistics**

Syntax Description

network	Clears counters for a FIB entry specified by network.
mask	(Optional) Clears counters for a FIB entry specified by network and mask.
*	Clears counters for all FIB entries.

Command Mode

EXEC

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

Related Commands

show adjacency
show ip cef

ip cef

To enable Cisco Express Forwarding (CEF) on the route processor card, use the **ip cef** global configuration command. To disable CEF, use the **no** form of this command.

ip cef [**distributed**] **switch**
no ip cef [**distributed**] **switch**

Syntax Description

distributed	(Optional) Enables distributed CEF (dCEF) operation. Distributes CEF information to line cards. Line cards perform express forwarding.

Defaults

On this platform...	The default is...
Cisco 7000 series equipped with RSP7000	CEF is not enabled.
Cisco 7200 series	CEF is not enabled.
Cisco 7500 series	CEF is not enabled.
Cisco 12000 series Gigabit Switch Router	Distributed CEF is enabled.

Command Mode

Global configuration

Usage Guidelines

This command first appeared in Cisco IOS Release 11.1 CC.

This command is not available on the Cisco 12000 series GSR because that router series operates only in distributed CEF mode.

CEF is advanced Layer 3 IP switching technology. CEF optimizes network performance and scalability for networks with dynamic, topologically dispersed traffic patterns, such as those associated with Web-based applications and interactive sessions.

Examples

The following example enables standard CEF operation:

```
ip cef switch
```

The following example enables dCEF operation:

```
ip cef distributed switch
```

Related Commands

ip route-cache cef

ip cef accounting

To enable network accounting of Cisco Express Forwarding (CEF), use the **ip cef accounting** global configuration command. To disable network accounting of CEF, use the **no** form of this command.

ip cef accounting [**per-prefix**] [**non-recursive**]
no ip cef accounting [**per-prefix**] [**non-recursive**]

Syntax Description

per-prefix	(Optional) Enables the collection of the number of packets and bytes express forwarded to a destination (or prefix).
non-recursive	(Optional) Enables accounting through non-recursive prefixes. For prefixes with directly connected next hops, enables the collection of the number of packets and bytes express forwarded through a prefix.

Default

Accounting is disabled by default.

Command Mode

Global configuration

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

You might want to collect statistics to better understand CEF patterns in your network.

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

You can then view the collected accounting information using the **show ip cef** command.

Related Commands

show ip cef

ip load-sharing

To enable load balancing for Cisco Express Forwarding (CEF), use the **ip load-sharing** interface configuration command. To disable load balancing for CEF, use the **no** form of this command.

ip load-sharing [**per-packet**] [**per-destination**]
no ip cef [**per-packet**]

Syntax Description

per-packet	(Optional) Enables per-packet load balancing on the interface.
per-destination	(Optional) Enables per-destination load balancing on the interface.

Default

Per-destination load balancing is enabled by default when you enable CEF.

Command Mode

Interface configuration

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

Per-packet load balancing allows the router to send data packets over successive equal-cost paths without regard to individual end hosts or user sessions. Path utilization is good, but packets destined for a given end host might take different paths and might arrive out of order.

Per-destination load balancing allows the router to use multiple, equal-cost paths to achieve load sharing. Packets for a given end host are guaranteed to take the same path, even if multiple, equal-cost paths are available. Traffic for different end hosts tend to take different paths.

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

Examples

The following example enables per-packet load balancing:

```
interface E0
 ip load-sharing per-packet
```

The following example enables per-destination load balancing:

```
interface E0
 ip load-sharing per-destination
```

ip route-cache cef

To enable Cisco Express Forwarding (CEF) operation on an interface after CEF operation has been disabled, use the **ip route-cache cef** interface configuration command. To disable CEF operation on an interface, use the **no** form of this command.

ip route-cache cef
no ip route-cache cef

Syntax Description

This command has no arguments or keywords.

Defaults

When standard CEF or dCEF operation is enabled globally, all interfaces that support CEF are enabled by default.

Command Mode

Interface configuration

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

Cisco Express Forwarding (CEF) is advanced Layer 3 switching technology for IP. CEF optimizes network performance and scalability for networks with dynamic, topologically dispersed traffic patterns, such as those associated with Web-based applications and interactive type sessions.

Because all interfaces that support CEF or dCEF are enabled by default when you enable standard CEF or dCEF operation globally, you use the **no** form of the command to turn off CEF operation on a particular interface.

You might want to disable CEF or dCEF on a particular interface because that interface is configured with a feature that CEF or dCEF does not support. For example, policy routing and CEF cannot be used together. You might want one interface to support policy routing while the other interfaces support CEF. In this case, you would turn on CEF globally, but turn off CEF on the interface configured for policy routing, enabling all but one interface to express forward.

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 reenable it, you can do so by using the **ip route-cache cef** command in interface configuration mode.

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

Examples

The following example enables CEF operation on the router (globally), but turns off CEF operation on Ethernet interface 0:

```
ip cef switch
interface e0
 no ip route-cache cef
```

The following example enables dCEF operation on the router (globally), but turns off CEF operation on Ethernet interface 0:

```
ip cef distributed switch
interface e0
 no ip route-cache cef
```

The following example reenables dCEF operation on Ethernet interface 0:

```
ip cef distributed switch
interface e0
 ip route-cache cef
```

Related Commands

interface
ip cef

show adjacency

To display Cisco Express Forwarding (CEF) adjacency table information, use the **show adjacency** EXEC command.

show adjacency [**detail**]

Syntax Description

detail	(Optional) Displays detailed adjacency information, including Layer 2 information.

Command Mode

EXEC

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

This command is available only on routers that have RP cards.

Sample Displays

The following is sample output from the **show adjacency detail** command.

```
Router# show adjacency detail

Protocol Interface                Address
IP       Ethernet1/0/0            9.2.61.1(7)
                                  0 packets, 0 bytes
                                  00107BC30D5C
                                  00500B32D8200800
                                  ARP        02:01:49
```

The encapsulation string 00107BC30D5C00500B32D8200800 is that of an adjacency used for traffic switched out of a router on an Ethernet link using Ethernet II encapsulation.

Table 2 describes the significant fields shown in the display.

Table 2 show adjacency detail Field Descriptions

Field	Description
Protocol	The routed protocol to which the adjacency is related.
Interface	The outgoing interface associated with the adjacency.
Address	The address can represent one of these addresses: Next Hop address Point-to-Point address The number (in parentheses) that follows this field indicates the number of internal references to the adjacency.
Source	The source where the adjacency was learned.

Table 2 show adjacency detail Field Descriptions (Continued)

Field	Description
Encapsulation string	The string which is prepended to a packet before the packet is transmitted.
Time stamp	The time left before the adjacency rolls out of the adjacency table. A packet must use the same next hop to the destination.

Related Commands

clear adjacency

show cef

To display which packets the line cards dropped or to display which packets were not express forwarded, use the **show cef** EXEC command.

show cef [**drop**] | [**not-cef-switched**]

Syntax Description

drop	(Optional) Displays which packets were dropped by each line card.
not-cef-switched	(Optional) Displays which packets were sent to a different switching path.

Command Mode

EXEC

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

This command is available only on routers that have RP cards.

A line card might drop packets due to encapsulation failure, no route information, or no adjacency information.

A packet is sent to a different switching path because Cisco Express Forwarding (CEF) does not support the encapsulation or feature, the packet is destined for the router, or the packet has IP options, such as time stamp and record route. IP options are process switched.

Sample Displays

The following is sample output from the **show cef drop** command:

```
Router# show cef drop

CEF Drop Statistics
Slot  Encap_fail  Unresolved Unsupported    No_route      No_adj   ChksumErr
RP             4          89           0           4           0           0
1              0           0           0           0           0           0
2              0           0           5           0           0           5
```

Table 3 describes the fields shown in the output.

Table 3 Show CEF Drop Field Descriptions

Field	Meaning
Slot	The slot number on which the packets were received.
Encap_fail	Indicates the number of packets dropped after the limit was reached for incomplete packets with no adjacency route.
Unresolved	Indicates the number of packets dropped because the route for the prefix was not resolved.

Field	Meaning
Unsupported	Indicates the number of packets received for which the adjacency route information was dropped due to unsupported features.
No_route	No route definition is included in the prefix table.
No_adj	The prefix is resolved, but the adjacent route is not indicated.
ChksumErr	Indicates the number of packets received with a checksum error.

The following is sample output from the **show cef not-cef-switched** command:

```
Router# show cef not-cef-switched

CEF Packets passed on to next switching layer
Slot No_adj No_encap Unsuppted Redirect Receive Bad_ttl Options Access RP
0    0      0        0          91584    0       0       0
1    0      0        0        0          0       0       0       0
2    0      0        0        0          0       0       0       0
```

Table 4 describes the fields shown in the output.

Table 4 Show CEF Not-CEF-Switched Field Descriptions

Field	Meaning
No_adj	Indicates the number of packets sent to the line card to ARP for the adjacent route.
No_encap	Number of encapsulated packets received.
Unsupported Redirect	Number of packets with unsupported features and redirected to another switching layer or location for processing.

Related Commands

show cef interface
show cef linecard

show cef interface

To display Cisco Express Forwarding (CEF) related interface information, use the **show cef interface** EXEC command.

show cef interface *type number* [**detail**]

Syntax Description

type number	Interface type and number about which to display CEF-related information.
detail	(Optional) Displays detailed CEF information for the specified interface type and number.

Command Mode

EXEC

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

This command is available on routers that have RP cards and line cards.

The **detail** command displays more CEF-related information for the specified interface.

You can use this command to show the CEF state on an individual interface.

Sample Displays

The following is sample output from the **show cef interface detail** command for Ethernet interface 0:

```
Router# show cef interface E0 detail

Ethernet1/0/0 is up (if_number 6)
Internet address is 172.19.177.20/24
ICMP redirects are always sent
Per-packet load balancing is disabled
Inbound access list is 10
Outbound access list is not set
Hardware idb is Ethernet1/0/0
Fast switching type 1, interface type 5
IP Distributed CEF switching enabled
IP Feature CEF switching turbo vector
Fast flags 0x4. ifindex 5(5)
Slot 1 Slot unit 0 VC -1
Hardware transmit queue ptr 0x48001A00 (0x48001A00) >- debugging purposes Transmit limit
accumulator 0x48001A02 (0x48001A02) IP MTU 1500
```

Table 5 describes the fields shown in the output.

Table 5 Show CEF Interface Detail Field Descriptions

Field	Meaning
interface type number is {up | down}	Indicates status of the interface.
Internet address	Internet address of the interface
ICMP packets are {always sent | never sent}	Indicates how packet forwarding is configured.
Per-packet load balancing	Status of load balancing in use on the interface (enabled or disabled).
Inbound access list {# | Not set}	Number of access lists defined for the interface.
Outbound access list	Number of access lists defined for the interface.
Hardware idb is *type number*	Interface type and number configured.
Fast switching type	Used for troubleshooting; indicates switching mode in use.
IP Distributed CEF switching {enabled | disabled}	Indicates the switching path used.
Slot *n* Slot unit *n*	The slot number.
Hardware transmit queue	Indicates the number of packets in the transmit queue.
Transmit limit accumulator	Indicates the maximum number of packets allowed in the transmit queue.
IP MTU	The value of the MTU size set on the interface.

Related Commands

show cef
show cef linecard

show cef linecard

To display Cisco Express Forwarding (CEF) related interface information by line card, use the **show cef linecard** EXEC command.

show cef linecard [*slot-number*] [**detail**]

Syntax Description

slot-number	(Optional) Slot number containing the line card about which to display CEF-related information. When you omit this argument, information about all line cards is displayed.
detail	(Optional) Displays detailed CEF information for the specified line card.

Command Mode

EXEC

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

This command is available only on routers that have RP cards.

When you omit the *slot-number* argument, information about all line cards is displayed. When you omit the *slot-number* argument and include the **detail** keyword, detailed information is displayed for all linecards. When you omit all keywords and arguments, the **show cef linecard** command displays important information about all line cards in table format.

Sample Displays

The following is sample output from the **show cef linecard** command. The command displays information for the line cards.

```
Router# show cef linecard

CEF table version 115705, 45877 routes
Slot CEF-ver MsgSent XdrSent Seq MaxSeq LowQ HighQ Flags
1        238     668    9641 616    616    0      0 up, sync
2        238     683   10782 619    629    0      0 up, sync
```

Table 6 describes the fields shown in the output.

Table 6 Show CEF Linecard Field Descriptions

Field	Meaning
CEF table version	The FIB table version.
XdrSent	IPC information elements (xdrs) packed into IPC messages sent from the RP to the line card.
MsgSent	Number of IPC messages sent.

Field	Meaning
Seq	Sequence number for the line card.
MaxSeq	Maximum sequence expected by the line card.
LowQ/HighQ	Number of xdr elements in LowQ and HighQ.
Flags	Indicates the status of the line card. Possible states are • up — Line card is up. • sync — Line card is in sync with main FIB. • repopulate — Repopulate FIB on line card. • reset — Line card FIB is reset. • reloading — Line card FIB is currently being reloaded. • disabled — Line card is disabled.

The following is sample output from the **show cef linecard detail** command for the line card in slot number 2:

```
Router# show cef linecard 2 detail

CEF line card slot number 2, status up, sync, disabled
line card CEF version number 238
Sequence number 616, Maximum sequence number expected 616
Send failed 0, Out Of Sequence 0
line card CEF reset 2, reloaded 2
92299/15/91 prefix/adjacency/interface elements queued
49641 elements packed in 668 messages(1341286 bytes) sent
0/0 xdr elements in LowQ/HighQ
Input packets 0, bytes 0<--- line card stats
Output packets 0, bytes 0, drops 0
```

Related Commands

show cef
show cef interface

show ip cef

To display entries in the FIB that are unresolved or to display a summary of the FIB, use this form of the **show ip cef** EXEC command:

show ip cef [**unresolved** | **summary**]

To display specific entries in the FIB based on IP address information, use this form of the **show ip cef** EXEC command:

show ip cef [*network* [*mask* [**longer-prefix**]]] [**detail**]

To display specific entries in the FIB based on interface information, use this form of the **show ip cef** EXEC command:

show ip cef [*type number*] [**detail**]

Syntax Description

unresolved	(Optional) Displays unresolved FIB entries.
summary	(Optional) Displays a summary of the FIB.
network	(Optional) Displays the FIB entry for the specified destination network.
mask	(Optional) Displays the FIB entry for the specified destination network and mask.
longer-prefix	(Optional) Displays FIB entries for all more specific destinations.
detail	(Optional) Displays detailed FIB entry information.
type number	(Optional) Interface type and number for which to display FIB entries.

Command Mode

EXEC

Usage Guidelines

This command first appeared to support the Cisco 12012 Gigabit Switch Router in Cisco IOS Release 11.2 GS, and first appeared with multiple platform support in Cisco IOS Release 11.1 CC.

The **show ip cef** command without any keywords or arguments shows a brief display of all FIB entries.

The **show ip cef detail** command shows detailed FIB entry information for all FIB entries.

The **show ip cef internal** command shows FIB load-sharing information for all FIB entries.

Sample Displays

The following is sample output from the **show ip cef unresolved** command:

```
Router# show ip cef unresolved

IP Distributed CEF with switching (Table Version 136632)
45776 routes, 13 unresolved routes (0 old, 13 new)
45776 leaves, 2868 nodes, 8441480 bytes, 136632 inserts, 90856 invalidations
1 load sharing elements, 208 bytes, 1 references
1 CEF resets, 1 revisions of existing leaves
refcounts: 527292 leaf, 465617 node

148.214.0.0/16, version 136622
0 packets, 0 bytes
  via 171.69.233.56, 0 dependencies, recursive
  unresolved
148.215.0.0/16, version 136623
0 packets, 0 bytes
  via 171.69.233.56, 0 dependencies, recursive
  unresolved
148.218.0.0/16, version 136624
0 packets, 0 bytes
```

The following is sample output from the **show ip cef summary** command:

```
Router# show ip cef summary

IP Distributed CEF with switching (Table Version 135165)
45788 routes, 0 reresolve, 4 unresolved routes (0 old, 4 new)
45788 leaves, 2868 nodes, 8442864 bytes, 135165 inserts, 89377 invalidations
0 load sharing elements, 0 bytes, 0 references
1 CEF resets, 0 revisions of existing leaves
refcounts: 527870 leaf, 466167 node
```

The following is sample output from the **show ip cef internal** command; it shows load-sharing details for multiple paths to a prefix:

```
Router# show ip cef 192.168.1.0 internal

192.168.1.0/24, version 135490, per-destination sharing 0 packets, 0 bytes

via 172.19.233.50, 0 dependencies, recursive<-- possible path 1 info
traffic share 1, current path
next hop 172.19.233.50, Ethernet0/0 via 172.19.233.50/32 valid adjacency
via 172.19.233.49, 0 dependencies, recursive<-- possible path 2 info
traffic share 1
next hop 172.19.233.49, Ethernet0/0 via 172.19.233.49/32 valid adjacency

0 packets, 0 bytes switched through the prefix Load distribution: 0 1 0 1 0 1 0 1 0 1 0
1 0 1 0 1 (refcount 1) ^
|.. how the load is distributed among the possible paths

Hash OK InterfaceAddressPackets
1Y Ethernet0/0172.19.233.500
2Y Ethernet0/0172.19.233.490
3Y Ethernet0/0172.19.233.500
4Y Ethernet0/0172.19.233.490
5Y Ethernet0/0172.19.233.500
6Y Ethernet0/0172.19.233.490
7Y Ethernet0/0172.19.233.500
8Y Ethernet0/0172.19.233.490
9Y Ethernet0/0172.19.233.500
10 Y Ethernet0/0172.19.233.490
11 Y Ethernet0/0172.19.233.500
```

```
12 Y Ethernet0/0172.19.233.490
13 Y Ethernet0/0172.19.233.500
14 Y Ethernet0/0172.19.233.490
15 Y Ethernet0/0172.19.233.500
16 Y Ethernet0/0172.19.233.490
```

The following is sample output from the **show ip cef detail** command for Ethernet interface 0. It shows all the prefixes resolving through adjacency pointing to next hop Ethernet interface 0/0 and next-hop interface IP address 172.19.233.33.

```
Router# show ip cef e0/0 172.19.233.33 detail
IP Distributed CEF with switching (Table Version 136808)
45800 routes, 8 unresolved routes (0 old, 8 new) 45800 leaves, 2868 nodes, 8444360 bytes,
136808 inserts, 91008 invalidations 1 load sharing elements, 208 bytes, 1 references 1
CEF resets, 1 revisions of existing leaves refcounts: 527343 leaf, 465638 node

172.19.233.33/32, version 7417, cached adjacency 172.19.233.33 0 packets, 0 bytes,
Adjacency-prefix
via 172.19.233.33, Ethernet0/0, 0 dependencies
next hop 172.19.233.33, Ethernet0/0
valid cached adjacency
```

Related Commands

show cef
show cef interface

Debug Command

The **debug ip cef command** is available for troubleshooting Cisco Express Forwarding (CEF) and distributed CEF.

debug ip cef

Use the **debug ip cef** EXEC command for troubleshooting CEF.

debug ip cef {**drops** [*access-list*] | **receive** [*access-list*] | **events** [*access-list*] | **prefix-ipc** [*access-list*] | **table** [*access-list*]}

and

debug ip cef {**ipc | interface-ipc**}

Syntax Description

drops	Records dropped packets.
access-list	Controls collection of debugging information from specified lists
receive	Records packets that are not switched using information from the FIB table, but are received and sent to the next switching layer
events	Records general CEF events.
prefix-ipc	Records updates related to IP prefix information. Possible updates include: • Debugging of IP routing updates in a line card • Reloading of a line card with a new table • Route update from the route processor to the line card has exceeded the maximum number of routes • Control messages related to FIB table prefixes
table	Produces a table showing events related to the FIB table. Possible types of events include: • Routing updates that populate the FIB table • Flushing of the FIB table • Adding or removing of entries to the FIB table • Table reloading process

ipc	Records information related to IPC in CEF. Possible types of events include: • Transmission status of IPC messages • Status of buffer space for ipc messages • IPC messages received out of sequence • Status of resequenced messages • Throttle requests sent from a line card to the route processor
interface- ipc	Records IPC updates related to interfaces. Possible reporting includes an interface coming up or going down, updates to fibhwidb, fibidb, and so forth.

What to Do Next

For information about the Cisco 12000 series router, refer to the *Cisco 12012 Gigabit Router Installation and Configuration Guide* (DOC-12012GSR-ICG=). For information about the Cisco 7500 series router, refer to the *Cisco 7500 Series Installation and Configuration Guide* (DOC-7500-ICG=).