



Cisco Express Forwarding Overview

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.

Procedures for configuring CEF or distributed CEF (dCEF) are provided in the “Configuring Cisco Express Forwarding” chapter later in this publication.

This chapter describes CEF. It contains the following sections:

- Benefits
- CEF Components
- Supported Media
- CEF Operation Modes
- Additional Capabilities

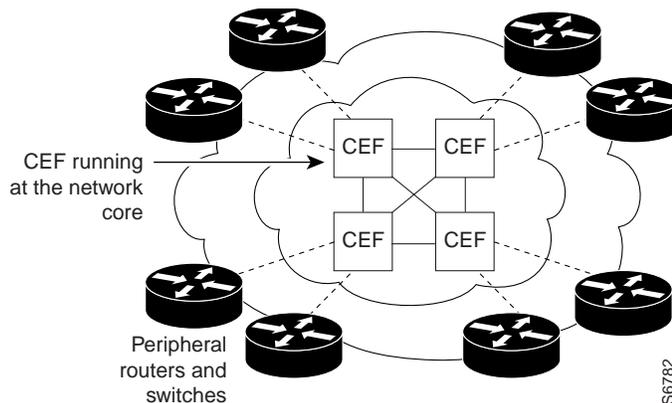
Benefits

CEF offers the following benefits:

- Improved performance—CEF is less CPU-intensive than fast 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 distributed CEF (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 Forwarding Information Base (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.

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 7 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 7 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 following:

- Forwarding Information Base
- Adjacency Tables

Forwarding Information Base

CEF uses a 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 switching paths such as fast switching and optimum switching.

Adjacency Tables

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 4.

Table 4 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 points 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.
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:

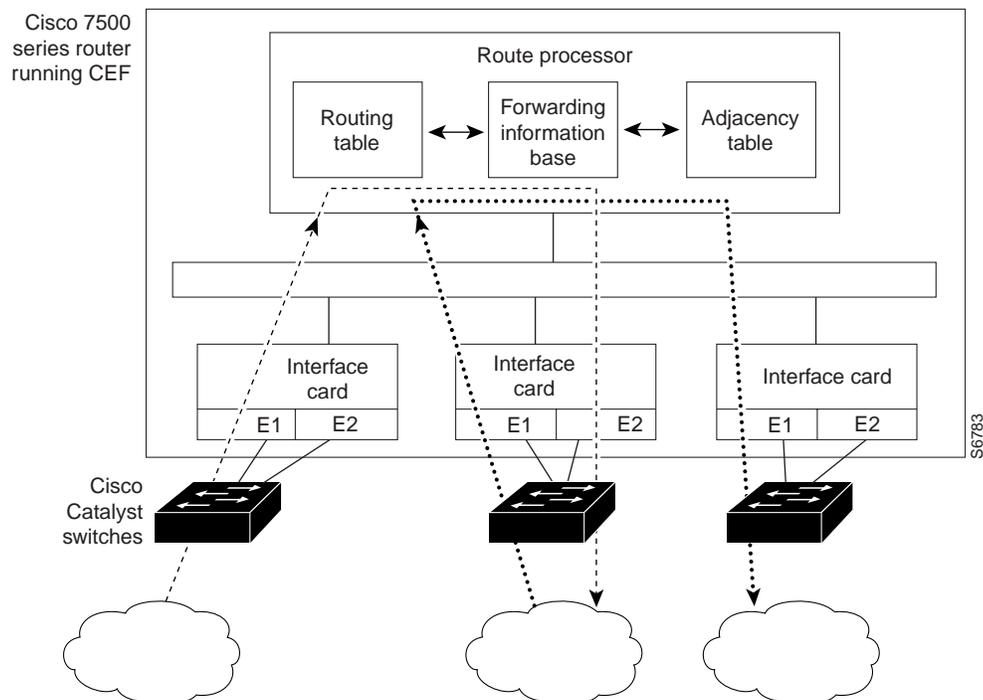
- Central CEF Mode
- Distributed CEF Mode

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 8 shows the relationship between the routing table, FIB, and adjacency table during CEF mode. The Cisco Catalyst switches forward traffic from workgroup LANs to a Cisco 7500 series router on the enterprise backbone running CEF. The route processor performs the express forwarding.

Figure 8 CEF Mode



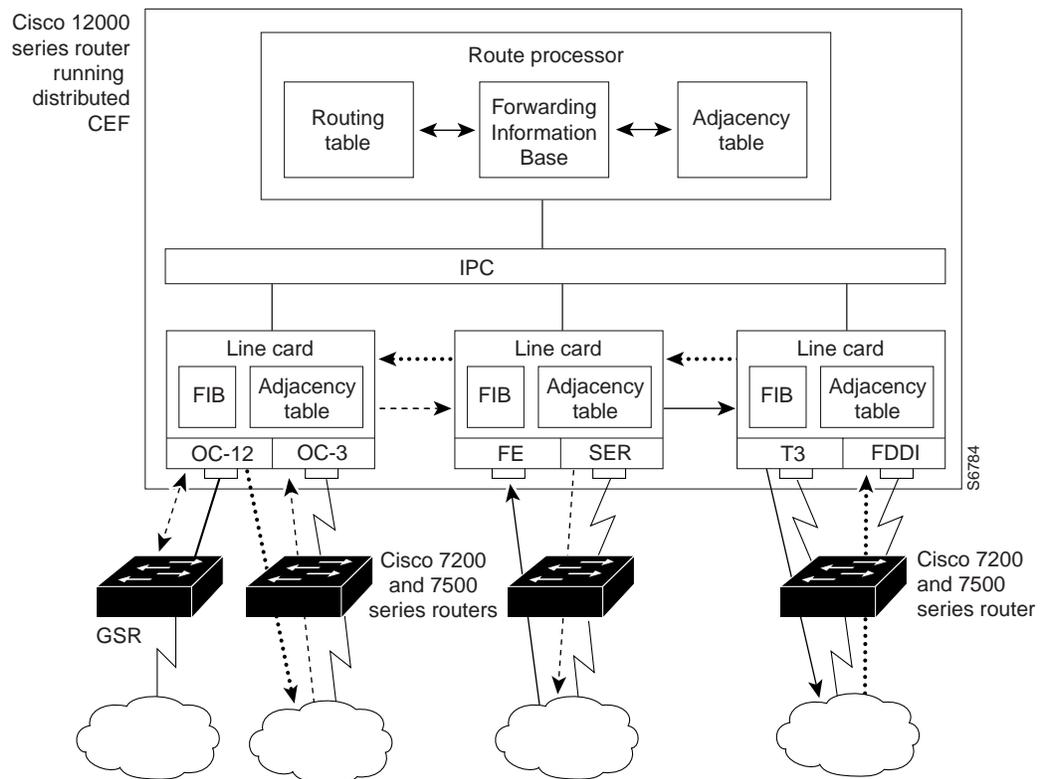
Distributed CEF Mode

When dCEF is enabled, line cards, such as VIP line cards or 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 9 shows the relationship between the route processor and line cards when dCEF mode is active.

Figure 9 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.



Note

The Cisco 12000 series Gigabit Switch Routers operate only dCEF mode; dCEF switching cannot be configured on the same VIP card as distributed fast switching, and dCEF is not supported on Cisco 7200 series routers.

Additional Capabilities

In addition to configuring CEF and dCEF, you can also configure the following 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-source host pair or per packet basis
- Network accounting to gather byte and packet statistics
- Distributed CEF switching across IP tunnels

For information on enabling these features, see the next chapter “Configuring Cisco Express Forwarding.”