



Cisco Express Forwarding Overview

First Published: May 2, 2005

Last Updated: July 31, 2009

This module contains an overview of the Cisco Express Forwarding feature. Cisco Express Forwarding is an advanced Layer 3 IP switching technology. It optimizes network performance and scalability for all kinds of networks: those that carry small amounts of traffic and those that carry large amounts of traffic in complex patterns, such as the Internet and networks characterized by intensive web-based applications or interactive sessions.

Finding Feature Information

For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the “[Feature Information for Cisco Express Forwarding Overview](#)” section on page 10.

Use Cisco Feature Navigator to find information about platform support and Cisco IOS XE software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

Contents

- [Information About Cisco Express Forwarding, page 2](#)
- [Additional References, page 8](#)
- [Feature Information for Cisco Express Forwarding Overview, page 10](#)
- [Glossary, page 11](#)

Information About Cisco Express Forwarding

Before using Cisco Express Forwarding or distributed Cisco Express Forwarding, you should understand the following:

- [Cisco Platform Support for Central Cisco Express Forwarding and Distributed Cisco Express Forwarding, page 2](#)
- [Cisco Express Forwarding Benefits: Improved Performance, Scalability, and Resilience, page 3](#)
- [Media Supported by Cisco Express Forwarding, page 4](#)
- [Main Components of Cisco Express Forwarding Operation, page 4](#)
- [FIB Overview, page 4](#)
- [Cisco Express Forwarding Adjacency Tables Overview, page 5](#)
- [Cisco Express Forwarding Operation, page 6](#)
- [Cisco Express Forwarding Features Enabled by Default, page 6](#)
- [Links for the Cisco Express Forwarding Features, page 7](#)

Cisco Platform Support for Central Cisco Express Forwarding and Distributed Cisco Express Forwarding

Cisco Express Forwarding is enabled by default on Cisco ASR 1000 Series Aggregation Service Routers running Cisco IOS XE software.

To find out if Cisco Express Forwarding is enabled on your platform, enter the **show ip cef** command. If Cisco Express Forwarding is enabled, you receive output that looks like this:

```
Router# show ip cef

Prefix          Next Hop           Interface
[...]
10.2.61.8/24    192.168.100.1     FastEthernet1/0/0
                  192.168.101.1     FastEthernet2/1/0
[...]
```

If Cisco Express Forwarding is not enabled on your platform, the output for the **show ip cef** command looks like this:

```
Router# show ip cef

%CEF not running
```

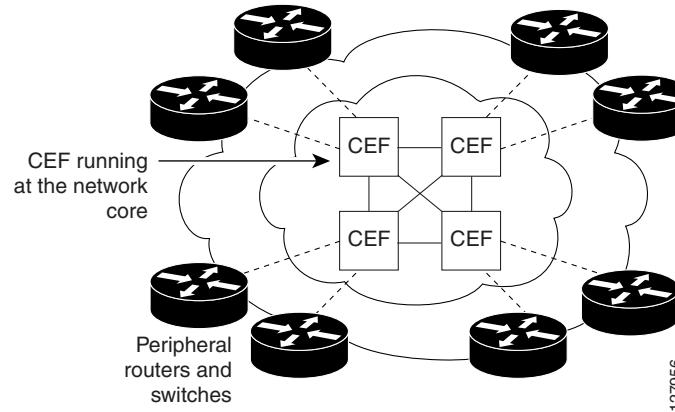
If Cisco Express Forwarding is not enabled on your platform, use the **ip cef** command to enable (central) Cisco Express Forwarding.

Cisco Express Forwarding Benefits: Improved Performance, Scalability, and Resilience

- Improved performance—Cisco Express Forwarding is less CPU-intensive than fast switching route caching. As a result, more CPU processing power can be dedicated to Layer 3 services such as quality of service (QoS) and encryption.
- Scalability—Cisco Express Forwarding offers full switching capacity at each line card when distributed Cisco Express Forwarding mode is active. Distributed Cisco Express Forwarding is a distributed switching mechanism that scales linearly with the number of interface cards and the bandwidth installed in the router.
- Resilience—Cisco Express Forwarding offers an unprecedented level of switching consistency and stability in large dynamic networks. In dynamic networks, fast-switched cache entries are frequently invalidated by routing changes. These changes can cause traffic to be process-switched through use of the routing table, rather than fast switched through use of the route cache. Because the forwarding information base (FIB) lookup table contains all known routes that exist in the routing table, it eliminates the need for route cache maintenance and the steps involved with fast-switch or process-switch forwarding. Cisco Express Forwarding can switch traffic more efficiently than typical demand caching schemes.

You can use Cisco Express Forwarding in any part of a network. For example, [Figure 1](#) shows Cisco Express Forwarding being run on routers at aggregation points at the core of a network where traffic levels are high and performance is critical.

Figure 1 Cisco Express Forwarding Example



Cisco Express Forwarding in platforms at the network core provides the performance and scalability that networks need to respond to continued growth and steadily increasing network traffic. Cisco Express Forwarding is a distributed switching mechanism that scales linearly with the number of interface cards and the bandwidth installed in the router.

Media Supported by Cisco Express Forwarding

Cisco Express Forwarding currently supports the following media:

- ATM/AAL5snap, ATM/AAL5mux, and ATM/AAL5nlpid
- Ethernet
- Frame Relay
- High-Level Data Link Control (HDLC)
- PPP
- Tunnels

Main Components of Cisco Express Forwarding Operation

Information conventionally stored in a route cache is stored in several data structures for Cisco Express Forwarding switching. The data structures provide optimized lookup for efficient packet forwarding. The two main components of Cisco Express Forwarding operation are the forwarding information base (FIB) and the adjacency tables.

The FIB is conceptually similar to a routing table or information base. A router uses this lookup table to make destination-based switching decisions during Cisco Express Forwarding operation. The FIB is updated when changes occur in the network and contains all routes known at the time. For more information, see the [“FIB Overview” section on page 4](#).

Adjacency tables maintain Layer 2 next-hop addresses for all FIB entries. For more information, see the [“Cisco Express Forwarding Adjacency Tables Overview” section on page 5](#).

This separation of the reachability information (in the Cisco Express Forwarding table) and the forwarding information (in the adjacency table), provides a number of benefits:

- The adjacency table can be built separately from the Cisco Express Forwarding table, allowing both to be built without any packets being processed.
- The MAC header rewrite used to forward a packet is not stored in cache entries, so changes in a MAC header rewrite string do not require invalidation of cache entries.

FIB Overview

Cisco Express Forwarding uses a FIB to make IP destination prefix-based switching decisions.

The FIB contains the prefixes from the IP routing table structured in a way that is optimized for forwarding. 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 the route cache maintenance that is associated with switching paths such as those used in fast switching and optimum switching.

Several paths can lead to a destination prefix. This occurs, for example, when a router is configured for simultaneous load balancing and redundancy. For each resolved path, the FIB contains a pointer for the adjacency corresponding to the next hop interface for that path.

Cisco Express Forwarding Adjacency Tables Overview

A node is said to be adjacent to another node if the node can be reached with a single hop across a link layer (Layer 2). Cisco Express Forwarding stores forwarding information (outbound interface and MAC header rewrite) for adjacent nodes in a data structure called the adjacency table. Cisco Express Forwarding uses adjacency tables to prepend Layer 2 addressing information to packets. The adjacency tables maintain Layer 2 next-hop addresses for all FIB entries.

The following sections provide additional information about adjacencies:

- [Adjacency Discovery, page 5](#)
- [Adjacency Types That Require Special Handling, page 5](#)
- [Unresolved Adjacency, page 6](#)

Adjacency Discovery

Each adjacency table is populated as adjacencies are discovered. Adjacencies are added to the table either through indirect manual configuration or dynamically—discovered through a mechanism like Address Resolution Protocol (ARP) or added through the use of a routing protocol, such as Border Gateway Protocol (BGP) or Open Shortest Path First (OSPF), which forms neighbor relationships. Each time an adjacency entry is created, a link-layer header for that adjacent node is computed and stored in the adjacency table.

The adjacency information is subsequently used for encapsulation during Cisco Express Forwarding switching of packets.

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. Prefixes requiring exception processing or special handling are cached with one of the special adjacencies listed in [Table 1](#).

Table 1 *Adjacency Types That Require Special Handling*

Packets of This Adjacency Type	Receive This Processing
Null adjacency	Packets destined for a Null0 interface are dropped. Null adjacency can be used as an effective form of access filtering.
Glean adjacency	When a router is connected to a multiaccess medium, 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. A glean adjacency entry indicates that a particular next hop should be directly connected, but there is no MAC header rewrite information available. When the router needs to forward packets to a specific host on a subnet, Cisco Express Forwarding requests an ARP entry for the specific prefix, ARP sends the MAC address, and the adjacency entry for the host is built.
Punt adjacency	The router forwards packets that require special handling or packets sent by features that are not yet supported in conjunction with Cisco Express Forwarding switching paths to the next higher switching level for handling.

Table 1 *Adjacency Types That Require Special Handling (continued)*

Packets of This Adjacency Type	Receive This Processing
Discard adjacency	The router discards the packets.
Drop adjacency	The router drops the packets.

Unresolved Adjacency

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

Cisco Express Forwarding Operation

You can use central Cisco Express Forwarding mode when line cards are not available for Cisco Express Forwarding switching, when you need to use features not compatible with distributed Cisco Express Forwarding switching, or when you are running on a nondistributed platform. When central Cisco Express Forwarding mode is enabled, the Cisco Express Forwarding FIB and adjacency tables reside on the RP, and the RP performs the express forwarding.

In routers where you can mix various types of cards in the same router, all cards might not support distributed Cisco Express Forwarding. When a line card that does not support distributed Cisco Express Forwarding receives a packet on one of these other routers, the line card forwards the packet to the next higher switching layer (the RP). This structure allows legacy interface processors to exist in the router with newer interface processors.



Note The Cisco ASR 1000 Series Aggregation Services Routers operate only in distributed Cisco Express Forwarding mode.

Cisco Express Forwarding Features Enabled by Default

- Per-destination load balancing and the universal load sharing algorithm (see the “[Configuring a Load-Balancing Scheme for Cisco Express Forwarding Traffic](#)” module)
- Distributed tunnel switching (see the “[Cisco Express Forwarding Distributed Tunnel Switching](#)” section on page 7)
- Multipoint generic routing encapsulation (GRE) tunnels (see the “[Cisco Express Forwarding-Switched Multipoint GRE Tunnels](#)” section on page 7)

The following sections describe Cisco Express Forwarding features enabled by default:

- [Cisco Express Forwarding Distributed Tunnel Switching, page 7](#)
- [Cisco Express Forwarding-Switched Multipoint GRE Tunnels, page 7](#)

Cisco Express Forwarding Distributed Tunnel Switching

Cisco Express Forwarding supports distributed tunnel switching, such as that made possible by GRE tunnels. Distributed tunnel switching is enabled automatically when you enable Cisco Express Forwarding or distributed Cisco Express Forwarding. You do not perform any additional tasks to enable distributed tunnel switching once you enable Cisco Express Forwarding or distributed Cisco Express Forwarding.

Cisco Express Forwarding-Switched Multipoint GRE Tunnels

The Cisco Express Forwarding-Switched Multipoint GRE Tunnels feature enables Cisco Express Forwarding switching of IP traffic to and from multipoint GRE tunnels. Traffic can be forwarded to a prefix through a tunnel destination when both the prefix and the tunnel destination are specified by the application. GRE creates a virtual point-to-point link to other routers at remote points over an IP internetwork. GRE can encapsulate a wide variety of protocol type packets. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows network expansion across a single-protocol backbone environment.

Links for the Cisco Express Forwarding Features

[Table 2](#) contains links to information about features that you can configure for use with Cisco Express Forwarding or distributed Cisco Express Forwarding operation.

Table 2 Features to Configure for Cisco Express Forwarding or Distributed Cisco Express Forwarding Operation

For Information on This Feature...	See the Following Document...
Configuring and verifying basic Cisco Express Forwarding operation	Configuring Basic Cisco Express Forwarding for Improved Performance, Scalability, and Resiliency in Dynamic Networks
Enabling or disabling Cisco Express Forwarding or distributed Cisco Express Forwarding switching and forwarding	Enabling or Disabling Cisco Express Forwarding or Distributed Cisco Express Forwarding to Customize Switching and Forwarding for Dynamic Networks
Changing your load-balancing scheme	Configuring a Load-Balancing Scheme for Cisco Express Forwarding Traffic
Refreshing or rebuilding adjacency or Cisco Express Forwarding tables	Configuring Epochs to Clear and Rebuild Cisco Express Forwarding and Adjacency Tables
Configuring Cisco Express Forwarding consistency checkers	Configuring Cisco Express Forwarding Consistency Checkers for Route Processors and Line Cards
Configuring network accounting for Cisco Express Forwarding	Configuring Cisco Express Forwarding Network Accounting
Customizing the display of Cisco Express Forwarding event trace messages	Customizing the Display of Cisco Express Forwarding Event Trace Messages

■ Additional References

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	<i>Cisco IOS Master Commands List, All Releases</i>
Commands for configuring and managing Cisco Express Forwarding	<i>Cisco IOS IP Switching Command Reference</i>
List of the features documented in the Cisco Express Forwarding modules	<i>Cisco Express Forwarding Features Roadmap</i>
Tasks for enabling or disabling Cisco Express Forwarding or distributed Cisco Express Forwarding	<i>Enabling or Disabling Cisco Express Forwarding or Distributed Cisco Express Forwarding to Customize Switching and Forwarding for Dynamic Networks</i>
Tasks for configuring a load-balancing scheme for Cisco Express Forwarding	<i>Configuring a Load-Balancing Scheme for Cisco Express Forwarding Traffic</i>
Tasks for configuring Cisco Express Forwarding consistency checkers	<i>Configuring Cisco Express Forwarding Consistency Checkers for Route Processors and Line Cards</i>
Tasks for configuring epochs for Cisco Express Forwarding tables	<i>Configuring Epochs to Clear and Rebuild Cisco Express Forwarding and Adjacency Tables</i>
Tasks for configuring and verifying Cisco Express Forwarding network accounting	<i>Configuring Cisco Express Forwarding Network Accounting</i>
Tasks for customizing the display of Cisco Express Forwarding event trace messages	<i>Customizing the Display of Cisco Express Forwarding Event Trace Messages</i>

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, Cisco IOS XE software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
RFC 1701	<i>Generic Route Encapsulation (GRE)</i>
RFC 2784	<i>Generic Routing Encapsulation (GRE)</i>
RFC 2890	<i>Key and Sequence Number Extensions to GRE</i>

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Cisco Express Forwarding Overview

Table 3 lists the features in this module.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which Cisco IOS XE software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.



Note

Table 3 lists only the Cisco IOS XE software release that introduced support for a given feature in a given Cisco IOS XE software release train. Unless noted otherwise, subsequent releases of that Cisco IOS XE software release train also support that feature.

Table 3 *Feature Information for Cisco Express Forwarding Overview*

Feature Name	Releases	Feature Configuration Information
Cisco Express Forwarding-Switched Multipoint GRE Tunnels	Cisco IOS XE Release 2.1	<p>This feature enables Cisco Express Forwarding switching of IP traffic to and from multipoint GRE tunnels. Prior to the introduction of this feature, only process switching was available for multipoint GRE tunnels.</p> <p>In Cisco IOS XE Release 2.1, this feature was introduced on the Cisco ASR 1000 Series Aggregation Services Routers.</p> <p>The following sections provide information about this feature:</p> <ul style="list-style-type: none"> • Cisco Express Forwarding-Switched Multipoint GRE Tunnels, page 7 <p>No commands were introduced or modified for this feature.</p>
CEF Support for IP Routing between IEEE 802.1Q VLANs	Cisco IOS XE Release 2.1	<p>In Cisco IOS XE Release 2.1, Cisco Express Forwarding support for IP routing between IEEE 802.1Q vLANs was introduced on Cisco ASR 1000 Series Aggregation Services Routers.</p> <p>No commands were introduced or modified for this feature.</p>

Glossary

adjacency—A relationship formed between selected neighboring routers and end nodes for the purpose of exchanging routing information. Adjacency is based upon the use of a common media segment by the routers and nodes involved.

Cisco Express Forwarding—A Layer 3 switching technology. Cisco Express Forwarding can also refer to central Cisco Express Forwarding mode, one of two modes of Cisco Express Forwarding operation. Cisco Express Forwarding enables a Route Processor to perform express forwarding. Distributed Cisco Express Forwarding is the other mode of Cisco Express Forwarding operation.

distributed Cisco Express Forwarding—A mode of Cisco Express Forwarding operation in which line cards maintain identical copies of the forwarding information base (FIB) and adjacency tables. The line cards perform the express forwarding between port adapters; this relieves the Route Processor of involvement in the switching operation.

FIB—forwarding information base. A component of Cisco Express Forwarding that is conceptually similar to a routing table or information base. The router uses the FIB lookup table to make destination-based switching decisions during Cisco Express Forwarding operation. The router maintains a mirror image of the forwarding information in an IP routing table.

GRE—generic routing encapsulation. A tunneling protocol developed by Cisco that enables encapsulation of a wide variety of protocol packet types inside IP tunnels, creating a virtual point-to-point link to Cisco routers at remote points over an IP internetwork. By connecting multiprotocol subnetworks in a single-protocol backbone environment, IP tunneling using GRE allows the expansion of a network across a single-protocol backbone environment.

IPC—interprocess communication. The mechanism that enables the distribution of Cisco Express Forwarding tables from the Route Processor (RP) to the line card when the router is operating in distributed Cisco Express Forwarding mode.

label disposition—The removal of Multiprotocol Label Switching (MPLS) headers at the edge of a network. In MPLS label disposition, packets arrive on a router as MPLS packets and, with the headers removed, are transmitted as IP packets.

label imposition—The action of putting a label on a packet.

LER—label edge router. A router that performs label imposition.

LFIB—label forwarding information base. The data structure used by switching functions to switch labeled packets.

LIB—label information base. A database used by a label switch router (LSR) to store labels learned from other LSRs, as well as labels assigned by the local LSR.

line card—A general term for an interface processor that can be used in various Cisco products.

LSP—label switched path. A sequence of hops (Router 0...Router n). A packet travels from R0 to Rn by means of label switching mechanisms. An LSP can be chosen dynamically, based on normal routing mechanisms, or it can be configured manually.

LSR—label switch router. A Layer 3 router that forwards a packet based on the value of a label encapsulated in the packet.

MPLS—Multiprotocol Label Switching. An emerging industry standard for the forwarding of packets along the normal routing paths (sometimes called MPLS hop-by-hop forwarding).

prefix—The network address portion of an IP address. A prefix is specified by a network and mask and is generally represented in the format network/mask. The mask indicates which bits are the network bits. For example, 1.0.0.0/16 means that the first 16 bits of the IP address are masked, making them the network bits. The remaining bits are the host bits. In this example, the network number is 10.0.

RIB—Routing Information Base. A central repository of routes that contains Layer 3 reachability information and destination IP addresses or prefixes. The RIB is also known as the routing table.

RP—Route Processor. The processor module in a router that contains the CPU, system software, and most of the memory components that are used in the router. It is sometimes called a supervisory processor.

VPN—Virtual Private Network. A router configuration that enables IP traffic to use tunneling to travel securely over a public TCP/IP network.

VRF—A Virtual Private Network (VPN) routing/forwarding instance. A VRF consists of an IP routing table, a derived forwarding table, a set of interfaces that use the forwarding table, and a set of rules and routing protocols that determine what goes into the forwarding table. In general, a VRF includes the routing information that defines a customer VPN site that is attached to a PE router.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2005–2009 Cisco Systems, Inc. All rights reserved.