



IP Switching Cisco Express Forwarding Configuration Guide, Cisco IOS XE Release 3SE (Catalyst 3850 Switches)

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: 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. (1110R)

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.

© 2013 Cisco Systems, Inc. All rights reserved.



CONTENTS

Cisco Express Forwarding	1
Finding Feature Information	1
Information About CEF	1
Cisco Platform Support for CEF and dCEF	2
CEF Benefits	2
Media Supported by CEF	3
Main Components of CEF	4
FIB Overview	4
CEF FIB and Load Balancing	4
CEF Adjacency Tables Overview	4
Adjacency Discovery	5
Adjacency Types That Require Special Handling	5
Unresolved Adjacency	6
Central CEF Mode Operation	6
Distributed CEF Mode Operation	6
CEF Features Enabled by Default	7
CEF Distributed Tunnel Switching	8
Links for the CEF Features	8
How to Configure CEF	8
Configuration Examples for CEF	9
Additional References	9
Feature Information for Cisco Express Forwarding	10



Cisco Express Forwarding

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, page 1](#)
- [Information About CEF, page 1](#)
- [How to Configure CEF, page 8](#)
- [Configuration Examples for CEF, page 9](#)
- [Additional References, page 9](#)
- [Feature Information for Cisco Express Forwarding, page 10](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and 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 table at the end of this module.

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

Information About CEF

- [Cisco Platform Support for CEF and dCEF, page 2](#)
- [CEF Benefits, page 2](#)
- [Media Supported by CEF, page 3](#)
- [Main Components of CEF, page 4](#)
- [FIB Overview, page 4](#)
- [CEF Adjacency Tables Overview, page 4](#)
- [Central CEF Mode Operation, page 6](#)
- [Distributed CEF Mode Operation, page 6](#)
- [CEF Features Enabled by Default, page 7](#)
- [Links for the CEF Features, page 8](#)

Cisco Platform Support for CEF and dCEF

Cisco Express Forwarding is enabled by default on most Cisco devices running Cisco software. When Cisco Express Forwarding is enabled on a device, the route processor (RP) performs the express forwarding.

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

```
Device# show ip cef
Prefix          Next Hop          Interface
[...]
```

Prefix	Next Hop	Interface
10.2.61.8/24	192.168.100.1	FastEthernet1/0/0
	192.168.101.1	FastEthernet6/1

```
[...]
```

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

```
Device# show ip cef
%CEF not running
```

When distributed Cisco Express Forwarding is enabled on your device, the line cards perform the express forwarding.

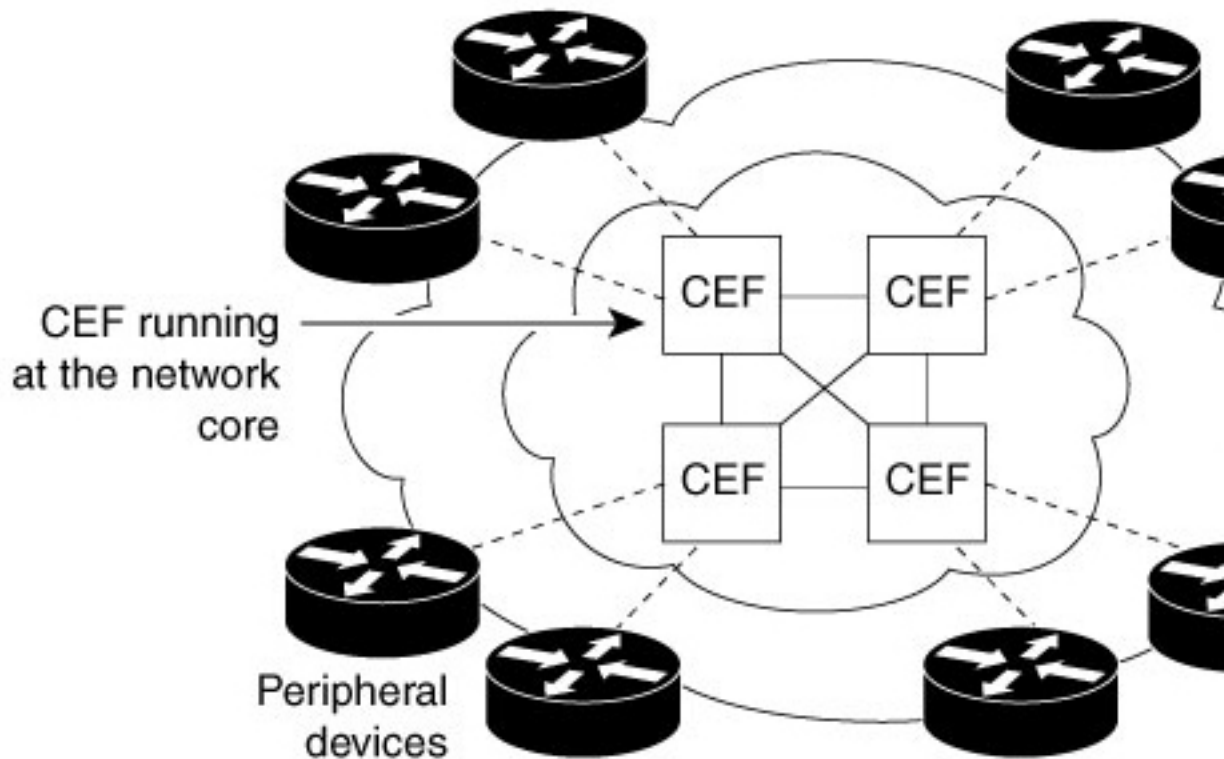
If Cisco Express Forwarding is not enabled on your device, use the **ip cef** command to enable Cisco Express Forwarding or use the **ip cef distributed** command to enable distributed Cisco Express Forwarding.

CEF Benefits

- 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 device.
- 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, the figure below shows Cisco Express Forwarding being run on devices at aggregation points at the core of a network where traffic levels are high and performance is critical.

Figure 1 Cisco Express Forwarding Example



Media Supported by CEF

Cisco Express Forwarding supports the following media:

- ATM/AAL5snap, ATM/AAL5mux, and ATM/AAL5nlpid
- Ethernet
- FDDI
- Frame Relay
- High-Level Data Link Control (HDLC)
- PPP
- Spatial Reuse Protocol (SRP)
- Token Ring
- Tunnels

Main Components of CEF

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.

Adjacency tables maintain Layer 2 next-hop addresses for all FIB entries. For more information, see the *CEF Adjacency Tables Overview* section.

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 process-switched.
- 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 validation of cache entries.

FIB Overview

Cisco Express Forwarding uses a forwarding information base (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.

- [CEF FIB and Load Balancing, page 4](#)

CEF FIB and Load Balancing

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.

CEF 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 through indirect manual configuration or dynamically--discovered through a mechanism like Address Resolution Protocol (ARP). Adjacencies can also be 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 the table below.

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 device is connected to a multiaccess medium, the FIB table on the device 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 device 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 device forwards packets requiring special handling or packets sent by features not yet supported in CEF switching paths to the next higher switching level for handling.
Discard adjacency	The device discards the packets.
Drop adjacency	The device drops the packets.

Unresolved Adjacency

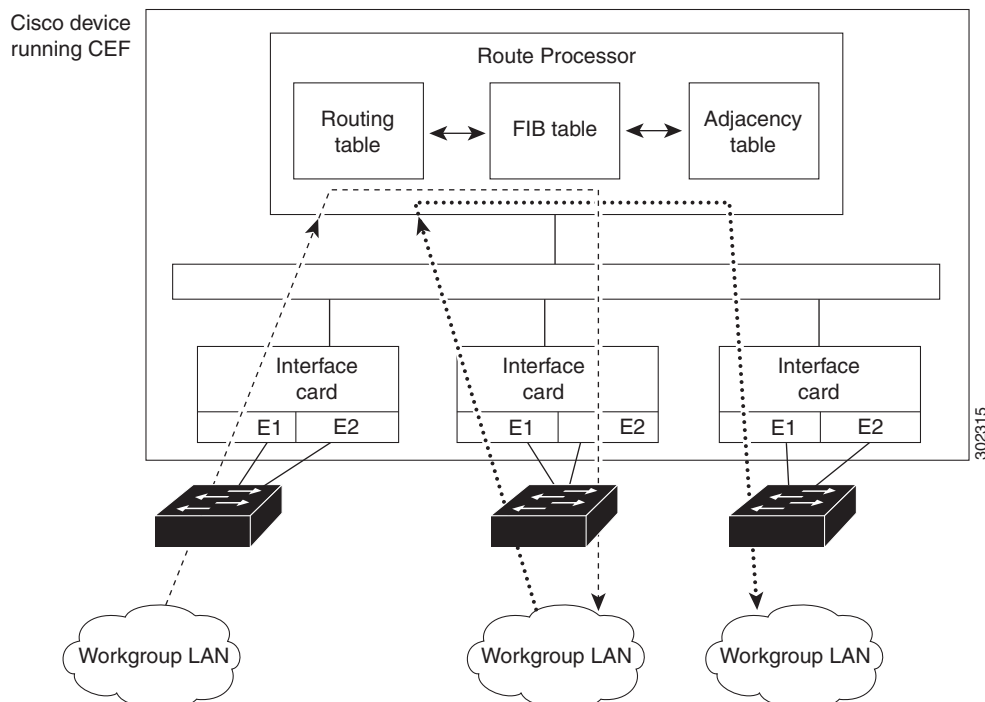
When a link-layer header is prepended to a packet, the FIB requires the prepended header to point to an adjacency corresponding to the next hop. If an adjacency was created by the FIB and not discovered through a mechanism such as ARP, the Layer 2 addressing information is not known and the adjacency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

Central CEF Mode 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 platform that is not a distributed 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.

The figure below shows the relationship between the routing table, the FIB, and the adjacency table during central Cisco Express Forwarding mode operation. Traffic is forwarded from workgroup LANs to a device on the enterprise backbone that is running central Cisco Express Forwarding. The RP performs the express forwarding.

Figure 2 Central Cisco Express Forwarding Mode Operation



Distributed CEF Mode Operation


For additional scalability, Cisco Express Forwarding runs in the distributed Cisco Express Forwarding form on certain devices by spreading processing tasks across two or more line cards. When distributed Cisco Express Forwarding mode is enabled, line cards maintain identical copies of the FIB and adjacency tables.

The line cards perform the express forwarding between port adapters, relieving the RP of involvement in the switching operation, thus also enhancing system performance.

Distributed Cisco Express Forwarding uses an interprocess communication (IPC) mechanism to ensure synchronization of FIB tables and adjacency tables on the RP and line cards.

The figure below shows the relationship between the RP and line cards when distributed Cisco Express Forwarding mode is active.

Figure 3 Distributed Cisco Express Forwarding Mode Operation



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

CEF Features Enabled by Default

- Per-destination load balancing and the universal load sharing algorithm
- Distributed tunnel switching
- Multipoint generic routing encapsulation (GRE) tunnels
- CEF distributed tunnel switching

- [CEF Distributed Tunnel Switching, page 8](#)

CEF Distributed Tunnel Switching

Cisco Express Forwarding supports distributed tunnel switching, such as that made possible by generic routing encapsulation (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.

Links for the CEF Features

The table below 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 CEF or dCEF Operation**

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

How to Configure CEF

There are no tasks for configuring Cisco Express Forwarding. Cisco Express Forwarding is enabled by default on most Cisco devices running Cisco software.

See the “Related Documents” section for links to configuration information for Cisco Express Forwarding features and services.

Configuration Examples for CEF

There are no configuration examples for the Cisco Express Forwarding.

See the “Related Documents” section for links to configuration information for Cisco Express Forwarding features and services.

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
IP switching commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples.	Cisco IOS IP Switching Command Reference
Tasks for verifying Cisco Express Forwarding information on your router	Configuring Basic Cisco Express Forwarding for Improved Performance, Scalability, and Resiliency in Dynamic Networks
Tasks for enabling or disabling Cisco Express Forwarding or distributed Cisco Express Forwarding	Enabling or Disabling Cisco Express Forwarding or Distributed Cisco Express Forwarding to Customize Switching and Forwarding for Dynamic Networks
Tasks for configuring a load-balancing scheme for Cisco Express Forwarding	Configuring a Load-Balancing Scheme for Cisco Express Forwarding Traffic
Tasks for configuring Cisco Express Forwarding consistency checkers	Configuring Cisco Express Forwarding Consistency Checkers for Route Processors and Line Cards
Tasks for configuring epochs for Cisco Express Forwarding tables	Configuring Epochs to Clear and Rebuild Cisco Express Forwarding and Adjacency Tables
Tasks for configuring and verifying Cisco Express Forwarding network accounting	Configuring Cisco Express Forwarding Network Accounting
Tasks for customizing the display of recorded Cisco Express Forwarding events	Customizing the Display of Recorded Cisco Express Forwarding Events
Verification steps for Cisco Express Forwarding switching	How to Verify Cisco Express Forwarding Switching
Troubleshooting tips for incomplete adjacencies	Troubleshooting Incomplete Adjacencies with CEF
Information about troubleshooting Cisco Express Forwarding routing loops and suboptimal routing	Troubleshooting Cisco Express Forwarding Routing Loops

Related Topic	Document Title
Causes of common Cisco Express Forwarding-related error messages on devices running distributed Cisco Express Forwarding switching and how to troubleshoot them	Troubleshooting Cisco Express Forwarding-Related Error Messages
Explanation of and troubleshooting information for the Cisco software implementation of Layer 3 load balancing across multiple parallel links when Cisco Express Forwarding is used	Troubleshooting Load Balancing Over Parallel Links Using Cisco Express Forwarding
QoS features that require Cisco Express Forwarding	When Is CEF Required for Quality of Service

RFCs

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

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

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

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

Table 3 **Feature Information for Cisco Express Forwarding**

Feature Name	Releases	Feature Configuration Information
CEF/dCEF - Cisco Express Forwarding	12.0 Cisco IOS XE Release 2.1	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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: 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. (1110R)

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.

