



Implementing Cisco Express Forwarding

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.



Note For complete descriptions of the CEF commands listed in this module, refer to the [Related Documents](#), on [page 13](#) section of this module.

Feature History for Implementing CEF

Release	Modification
Release 5.0.0	This feature was introduced.
Release 5.0.1	Support for Loose and Strict uRPF was added.

- [Prerequisites for Implementing Cisco Express Forwarding](#), on page 1
- [Information About Implementing Cisco Express Forwarding Software](#), on page 2
- [How to Implement CEF](#), on page 5
- [Configuration Examples for Implementing CEF on Routers Software](#), on page 11
- [Additional References](#), on page 12

Prerequisites for Implementing Cisco Express Forwarding

The following prerequisites are required to implement Cisco Express Forwarding:

- You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Information About Implementing Cisco Express Forwarding Software

To implement Cisco Express Forwarding features in this document you must understand the following concepts:

Key Features Supported in the Cisco Express Forwarding Implementation

The following features are supported for CEF on Cisco IOS XR software:

- Multipath support
- High availability features such as packaging, restartability, and Out of Resource (OOR) handling
- OSPFv2 SPF prefix prioritization
- BGP attributes download

Benefits of CEF

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.
- Resilience—CEF offers an unprecedented level of switching consistency and stability in large dynamic networks. In dynamic networks, fast-switched 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 or process-switch forwarding scenario. CEF can switch traffic more efficiently than typical demand caching schemes.

CEF Components

Cisco IOS XR software CEF always operates in CEF mode with two distinct components: a Forwarding Information Base (FIB) database and adjacency table—a protocol-independent adjacency information base (AIB).

CEF is a primary IP packet-forwarding database for Cisco IOS XR software. CEF is responsible for the following functions:

- Software switching path
- Maintaining forwarding table and adjacency tables (which are maintained by the AIB) for software and hardware forwarding engines

The following CEF forwarding tables are maintained in Cisco IOS XR software:

- IPv4 CEF database
- IPv6 CEF database
- MPLS LFD database
- Multicast Forwarding Table (MFD)

The protocol-dependent FIB process maintains the forwarding tables for IPv4 and IPv6 unicast in the () and each MSC.

The FIB on each node processes Routing Information Base (RIB) updates, performing route resolution and maintaining FIB tables independently in the and each MSC. FIB tables on each node can be slightly different. Adjacency FIB entries are maintained only on a local node, and adjacency entries linked to FIB entries could be different.

Reverse Path Forwarding (Strict and Loose)

Unicast IPv4 and IPv6 Reverse Path Forwarding (uRPF), both strict and loose modes, help mitigate problems caused by the introduction of malformed or spoofed IP source addresses into a network by discarding IP packets that lack a verifiable IP source address. Unicast RPF does this by doing a reverse lookup in the CEF table. Therefore, Unicast Reverse Path Forwarding is possible only if CEF is enabled on the router.



Note Unicast RPF allows packets with 0.0.0.0 source addresses and 255.255.255.255 destination addresses to pass so that Bootstrap Protocol and Dynamic Host Configuration Protocol (DHCP) will function properly.

When strict uRPF is enabled, the source address of the packet is checked in the FIB. If the packet is received on the same interface that would be used to forward the traffic to the source of the packet, the packet passes the check and is further processed; otherwise, it is dropped. Strict uRPF should only be applied where there is natural or configured symmetry. Because internal interfaces are likely to have routing asymmetry, that is, multiple routes to the source of a packet, strict uRPF should not be implemented on interfaces that are internal to the network.



Note The behavior of strict RPF varies slightly by platform, number of recursion levels, and number of paths in Equal-Cost Multipath (ECMP) scenarios. A platform may switch to loose RPF check for some or all prefixes, even though strict RPF is configured.

When loose uRPF is enabled, the source address of the packet is checked in the FIB. If it exists and matches a valid forwarding entry, the packet passes the check and is further processed; otherwise, it is dropped.

Loose and strict uRPF supports two options: **allow self-ping** and **allow default**. The **self-ping** option allows the source of the packet to ping itself. The **allow default** option allows the lookup result to match a default routing entry. When the **allow default** option is enabled with the strict mode of the uRPF, the packet is processed further only if it arrived through the default interface.

Route Processor Management Ethernet Forwarding

Forwarding from the MSC interface to the RP Management Ethernet is disabled by default. The **rp mgmtethernet forwarding** command is used to enable forwarding from the MSC interface to RP Management Ethernet.

Forwarding from the RP Management Ethernet to the MSC interface, and from the RP Management Ethernet to RP Management Ethernet, is enabled by default.

Per-Flow Load Balancing

Load balancing describes the functionality in a router that distributes packets across multiple links based on Layer 3 (network layer) and Layer 4 (transport layer) routing information. If the router discovers multiple paths to a destination, the routing table is updated with multiple entries for that destination.

Per-flow load balancing performs these functions:

- Incoming data traffic is evenly distributed over multiple equal-cost connections within a bundle interface.
- Layer 2 bundle and Layer 3 (network layer) load balancing decisions are taken on IPv4, IPv6, which are supported for the 7-tuple hash algorithm.
- A 7-tuple hash algorithm provides more granular load balancing than the existing 3-tuple hash algorithm.
- The same hash algorithm (3-tuple or 7-tuple) is used for load balancing over multiple equal-cost Layer 3 (network layer) paths. The Layer 3 (network layer) path is on a physical interface or on a bundle interface. In addition, load balancing over member links can occur within a Layer 2 bundle interface.
- The **cef load-balancing fields** command allows you to select either the 3-tuple hash algorithm (default) or the 7-tuple hash algorithm.

Layer 3 (Network Layer) Routing Information

The 3-tuple load-balance hash calculation contains these Layer 3 (Network Layer) inputs:

- Source IP address
- Destination IP address
- Router ID

The 7-tuple load-balance hash calculation contains 3-tuple inputs and these additional following Layer 4 (Transport Layer) inputs:

Layer 4 (Transport Layer) Routing Information

The 5-tuple load-balance hash calculation contains 3-tuple inputs and these additional following Layer 4 (Transport Layer) inputs:

- Source port
- Destination port
- Protocol



Note In load-balancing scenarios, a line card may not use all output paths downloaded from routing protocols. This behavior varies with platform, number of recursion levels, and the fact whether MPLS is involved, or not.

BGP Attributes Download

The BGP Attributes Download feature enables you to display the installed BGP attributes in CEF. Configure the **show cef bgp-attribute** command to display the installed BGP attributes in CEF. You can use the **show cef bgp-attribute attribute-id** command and the **show cef bgp-attribute local-attribute-id** command to look at specific BGP attributes by attribute ID and local attribute ID.

Verification

```
Router# show cef bgp-attribute
Wed Aug 21 14:05:51.772 UTC
```

```
VRF: default
```

```
Table ID: 0xe0000000. Total number of entries: 1
OOR state: GREEN. Number of OOR attributes: 0
```

```
BGP Attribute ID: 0x6, Local Attribute ID: 0x1
  Aspath      : 2
  Community   :
  Origin AS   : 2
  Next Hop AS : 2
```

How to Implement CEF

This section contains instructions for the following tasks:

Verifying CEF

This task allows you to verify CEF.

SUMMARY STEPS

1. **show cef {ipv4 | ipv6}**
2. **show cef {ipv4 | ipv6} summary**
3. **show cef {ipv4 | ipv6} detail**
4. **show adjacency detail**

DETAILED STEPS

	Command or Action	Purpose
Step 1	show cef {ipv4 ipv6} Example:	Displays the IPv4 or IPv6 CEF table. The next hop and forwarding interface are displayed for each prefix.

	Command or Action	Purpose
	RP/0/RP0/CPU0:router# show cef ipv4	Note The output of the show cef command varies by location.
Step 2	show cef {ipv4 ipv6} summary Example: RP/0/RP0/CPU0:router# show cef ipv4 summary	Displays a summary of the IPv4 or IPv6 CEF table.
Step 3	show cef {ipv4 ipv6} detail Example: RP/0/RP0/CPU0:router# show cef ipv4 detail	Displays detailed IPv4 or IPv6 CEF table information.
Step 4	show adjacency detail Example: RP/0/RP0/CPU0:router# show adjacency detail	Displays detailed adjacency information, including Layer 2 information for each interface. Note The output of the show adjacency command varies by location.

Configuring a Route Purge Delay

This task allows you to configure a route purge delay. A purge delay purges routes when the RIB or other related process experiences a failure.

SUMMARY STEPS

1. **configure**
2. **cef purge-delay** *seconds*
3. **commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	
Step 2	cef purge-delay <i>seconds</i> Example: RP/0/RP0/CPU0:router(config)# cef purge-delay 180	Configures a delay in purging routes when the Routing Information Base (RIB) or other related processes experience a failure.
Step 3	commit	

Configuring Unicast RPF Checking

This task allows you to configure unicast Reverse Path Forwarding (uRPF) checking. Unicast RPF checking allows you to mitigate problems caused by malformed or forged (spoofed) IP source addresses that pass

through a router. Malformed or forged source addresses can indicate denial-of-service (DoS) attacks based on source IP address spoofing.

SUMMARY STEPS

1. **configure**
2. **interface** *type interface-path-id*
3. **ipv4 verify unicast source reachable-via** {any | rx} [**allow-default**] [**allow-self-ping**]
4. **commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	
Step 2	interface <i>type interface-path-id</i> Example: RP/0/RP0/CPU0:router(config)# interface TenGigE 0/1/0/0	Enters interface configuration mode.
Step 3	ipv4 verify unicast source reachable-via {any rx} [allow-default] [allow-self-ping] Example: RP/0/RP0/CPU0:router(config-if)# ipv4 verify unicast source reachable-via rx	Enables IPv4 uRPF checking. <ul style="list-style-type: none"> • The rx keyword enables strict unicast RPF checking. If strict unicast RPF is enabled, a packet is not forwarded unless its source prefix exists in the routing table and the output interface matches the interface on which the packet was received. • The allow-default keyword enables the matching of default routes. This option applies to both loose and strict RPF. • The allow-self-ping keyword enables the router to ping out an interface. This option applies to both loose and strict RPF.
Step 4	commit	

Configuring Modular Services Card-to-Route Processor Management Ethernet Interface Switching

This task allows you to enable MSC-to-RP management Ethernet interface switching.

SUMMARY STEPS

1. **configure**
2. **rp mgmtethernet forwarding**
3. **commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>configure</code>	
Step 2	rp mgmtethernet forwarding Example: <pre>RP/0/RP0/CPU0:router(config)# rp mgmtethernet forwarding</pre>	Enables switching from the MSC to the route processor Management Ethernet interfaces.
Step 3	<code>commit</code>	

Configuring Per-Flow Load Balancing

This section describes the following tasks to configure per-flow load balancing:

Configuring a 7-Tuple Hash Algorithm

This task allows you to configure per-flow load balancing for a 7-tuple hash algorithm.

SUMMARY STEPS

1. `configure`
2. `cef load-balancing fields {L3 | L4}`
3. `commit`
4. `show cef {ipv4 | ipv6} summary [location node-id]`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>configure</code>	
Step 2	cef load-balancing fields {L3 L4} Example: <pre>RP/0/RP0/CPU0:router(config)# cef load-balancing fields L4</pre>	<p>Configures the hashing algorithm that is used for load balancing during forwarding. The example shows that the L4 field is selected.</p> <ul style="list-style-type: none"> • Use the L3 keyword to specify the Layer 3 load-balancing for the hash <p>Since L3 is configured as the default value, you do not need to use the cef load-balancing fields command unless you want to configure Layer 4.</p> <ul style="list-style-type: none"> • Use the L4 keyword to specify the Layer 3 and Layer 4 load-balancing for the hash algorithm. <p>For a list of the inputs for Layer 3 and Layer 4, see Per-Flow Load Balancing, on page 4.</p>
Step 3	<code>commit</code>	

	Command or Action	Purpose
Step 4	show cef {ipv4 ipv6} summary [location node-id] Example: <pre>RP/0/RP0/CPU0:router# show cef ipv4 summary</pre>	Displays the load balancing field for the IPv4 or IPv6 CEF table. <ul style="list-style-type: none"> (Optional) Use the location keyword display a summary of the IPv4 CEF table for the designated node. The <i>node-id</i> argument is entered in the <i>rack/slot/module</i> notation

Verifying the CEF Exact Route with 7-Tuple Parameters

The following 7-tuple parameters are specified to obtain the CEF exact route for both IPv4 and IPv6:

- Source address
- Destination address
- Source port and range of destination ports
- Protocol
- Ingress interface
- Router ID

To display the path an MPLS flow would take, use the

SUMMARY STEPS

1. Configure parallel interfaces between back-to-back routers.
2. Create route traffic streams so that there is a stream placed onto each configured interface.
3. Use the **show cef ipv4 exact-route** command in XR EXEC mode to verify that the interface selected for load balancing matches with the output from this command. The following example shows the exact route for the Layer 4 information:
4. Configure Equal Cost Multipath Protocol (ECMP) interfaces, for example, between back-to-back routers.
5. Create route traffic streams so that there is a stream placed onto each configured interface.
6. Use the **show cef ipv6 exact-route** command in XR EXEC mode to verify that the interface selected for load balancing matches with the output from this command. The following example shows the exact route for the Layer 4 information:

DETAILED STEPS

-
- Step 1** Configure parallel interfaces between back-to-back routers.
- Step 2** Create route traffic streams so that there is a stream placed onto each configured interface.
- Step 3** Use the **show cef ipv4 exact-route** command in XR EXEC mode to verify that the interface selected for load balancing matches with the output from this command. The following example shows the exact route for the Layer 4 information:

Example:

```
RP/0/RP0/CPU0:router# show cef ipv4 exact-route 20 .6.1.9 22.6.1.9 protocol udp source-port 1
destination-port 1 ingress-interface HundredGigE 0/1/0/4
```

```
22.6.1.9/32 version 0, internal 0x40040001 (0x78439fd0) [3], 0x0 (0x78aaf928), 0x4400 (0x78ed62d0)
  remote adjacency to HundredGigE0/1/4/4 Prefix Len 32, traffic index 0, precedence routine (0)
  via HundredGigE0/1/4/4
```

To verify the IPv6 7-tuple parameters, perform the following steps:

- Step 4** Configure Equal Cost Multipath Protocol (ECMP) interfaces, for example, between back-to-back routers.
- Step 5** Create route traffic streams so that there is a stream placed onto each configured interface.
- Step 6** Use the **show cef ipv6 exact-route** command in XR EXEC mode to verify that the interface selected for load balancing matches with the output from this command. The following example shows the exact route for the Layer 4 information:

Example:

```
RP/0/RP0/CPU0:router# show cef ipv6 exact-route 20:6:1::9 22:6:1::9 protocol udp source-port 1
destination-port 1 ingress-interface HundredGigE 0/1/0/4
```

```
22:6:1::/64, version 0, internal 0x40000001 (0x7846c048) [3], 0x0 (0x78aea3d0), 0x0 (0x0) remote
adjacency to HundredGigE0/1/4/4 Prefix Len 64, traffic index 0, precedence routine (0)
  via HundredGigE0/1/4/4
```

Configuring BGP Attributes Download

This task allows you to configure the BGP Attributes Download feature.

Configuring BGP Attributes Download

SUMMARY STEPS

1. **configure**
2. **cef bgp attribute** {*attribute-id* | *local-attribute-id*}
3. **commit**

DETAILED STEPS

	Command or Action	Purpose
Step 1	configure	
Step 2	cef bgp attribute { <i>attribute-id</i> <i>local-attribute-id</i> } Example: RP/0/RP0/CPU0:router(config)# cef bgp attribute	Configures a CEF BGP attribute.
Step 3	commit	

Configuration Examples for Implementing CEF on Routers Software

This section provides the following configuration examples:

Configuring Unicast RPF Checking: Example

The following example shows how to configure unicast RPF checking:

```
configure
interface TenGigE 0/0/0/1
ipv4 verify unicast source reachable-via rx
end
```

Configuring the Switching of Modular Services Card to Management Ethernet Interfaces on the Route Processor: Example

The following example shows how to configure the switching of the MSC to Management Ethernet interfaces on the route processor:

```
configure
rp mgmtethernet forwarding
end
```

Configuring Per-Flow Load Balancing: Example

The following examples show how to configure Layer 3 and Layer 4 load-balancing for the hash algorithm from the **cef load-balancing fields** command, and how to verify summary information for the CEF table from the **show cef summary** command:

Configuring Layer 3 load-balancing

```
configure
  cef load-balancing fields L3
end
!
show cef summary
Router ID is 10.6.6.6

IP CEF with switching (Table Version 0) for node0_RP0_CPU0

Load balancing: L3
Tableid 0xe0000000 (0x9cbb51b0),  Flags 0x2031
  Refcount 577
  300 routes, 0 protected, 0 reresolve, 0 unresolved (0 old, 0 new), 21600 bytes
  212 load sharing elements, 62576 bytes, 324 references
  19 shared load sharing elements, 5388 bytes
  193 exclusive load sharing elements, 57188 bytes
  0 route delete cache elements
  622 local route bufs received, 1 remote route bufs received,  0 mix bufs received
  176 local routes, 0 remote routes
```

```

4096 total local route updates processed
0 total remote route updates processed
0 pkts pre-routed to cust card

0 pkts received from core card
0 CEF route update drops, 96 revisions of existing leaves
0 CEF route update drops due to version mis-match
Resolution Timer: 15s
0 prefixes modified in place
0 deleted stale prefixes
82 prefixes with label imposition, 107 prefixes with label information
95 next hops
0 incomplete next hops

0 PD backwalks on LDIs with backup path

```

Configuring Layer 4 load-balancing

Router ID is 1.10.10.10

IP CEF with switching (Table Version 0) for node0_RP0_CPU0

```

Load balancing: L4
Tableid 0xe0000000 (0x89bba258), Flags 0x2031
RefCount 16
5 routes, 0 protected, 0 reresolve, 0 unresolved (0 old, 0 new), 680 bytes
5 load sharing elements, 1860 bytes, 0 references
0 shared load sharing elements, 0 bytes
5 exclusive load sharing elements, 1860 bytes
0 route delete cache elements
0 local route bufs received, 1 remote route bufs received, 0 mix bufs received
0 local routes, 0 remote routes
0 total local route updates processed
0 total remote route updates processed
0 pkts pre-routed to cust card
0 pkts received from core card
0 CEF route update drops, 0 revisions of existing leaves
0 CEF route update drops due to version mis-match
Resolution Timer: 15s
0 prefixes modified in place
0 deleted stale prefixes
0 prefixes with label imposition, 0 prefixes with label information
0 next hops
0 incomplete next hops

0 PD backwalks on LDIs with backup path

```

Configuring BGP Attributes Download: Example

The following example shows how to configure the BGP Attributes Download feature:

```

router configure
show cef bgp attribute {attribute-id| local-attribute-id}

```

Additional References

The following sections provide references related to implementing CEF.

Related Documents

Related Topic	Document Title
CEF commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Cisco Express Forwarding Commands</i> module in <i>IP Addresses and Services Command Reference for Cisco NCS 6000 Series Routers</i>
BGP commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>BGP Commands</i> module in the <i>Routing Command Reference for Cisco NCS 6000 Series Routers</i>
Link Bundling Commands: complete command syntax, command modes, command history, defaults, usage guidelines, and examples	<i>Link Bundling Commands</i> module in the <i>Interface and Hardware Component Command Reference for the Cisco NCS 6000 Series Routers</i>

Standards

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

MIBs

MIBs	MIBs Link
—	To locate and download MIBs, use the Cisco MIB Locator found at the following URL and choose a platform under the Cisco Access Products menu: https://mibs.cloudapps.cisco.com/ITDIT/MIBS/servlet/index

RFCs

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

Technical Assistance

Description	Link
The Cisco Technical Support website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

