



Implementing Cisco Express Forwarding

- [Implementing Cisco Express Forwarding, on page 1](#)

Implementing Cisco Express Forwarding

Cisco Express Forwarding (CEF) is an 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. CEF is an inherent feature and the users need not perform any configuration to enable it. If required, the users can change the default route purge delay and static routes. Cisco NCS 5500 Series Routers supports only single stage forwarding.

Components

Cisco IOS XR software CEF always operates in CEF mode with two distinct components:

- Forwarding Information Base (FIB) database: The protocol-dependent FIB process maintains the forwarding tables for IPv4 and IPv6 unicast in the route processor and line card (LC). The FIB on each node processes Routing Information Base (RIB) updates, performing route resolution and maintaining FIB tables independently in the route processor and line card (LC). FIB tables on each node can be slightly different.
- 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 features are supported for CEF on Cisco IOS XR software:

- Bundle interface support
- Multipath support
- Route consistency
- High availability features such as packaging, restartability, and Out of Resource (OOR) handling
- OSPFv2 SPF prefix prioritization

- BGP attributes download

CEF 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.

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

- IPv4 CEF database—Stores IPv4 Unicast routes for forwarding IPv4 unicast packets
- IPv6 CEF database—Stores IPv6 Unicast routes for forwarding IPv6 unicast packets
- MPLS LFD database—Stores MPLS Label table for forwarding MPLS packets

Verifying CEF

To view the details of the IPv4 or IPv6 CEF tables, use the following commands:

- `show cef {ipv4 address | ipv6 address} hardware egress`

Displays the IPv4 or IPv6 CEF table. The next hop and forwarding interface are displayed for each prefix. The output of the **show cef** command varies by location.

```
Router# show cef 203.0.1.2 hardware egress
 203.0.1.2/32, version 0, internal 0x1020001 0x0 (ptr 0x8d7db7f0) [1], 0x0 (0x8daeedf0),
0x0 (0x0)
Updated Nov 20 13:33:23.557
local adjacency 203.0.1.2
Prefix Len 32, traffic index 0, Adjacency-prefix, precedence n/a, priority 15
  via 203.0.1.2/32, HundredGigE0/0/0/9, 3 dependencies, weight 0, class 0 [flags 0x0]
  path-idx 0 NHID 0x0 [0x8cfc81a0 0x0]
  next hop 203.0.1.2/32
  local adjacency
```

- `show cef {ipv4 | ipv6} summary`

Displays a summary of the IPv4 or IPv6 CEF table.

```
Router#show cef ipv4 summary
Fri Nov 20 13:50:45.239 UTC

Router ID is 216.1.1.1

IP CEF with switching (Table Version 0) for node0_RP0_CPU0

Load balancing: L4
Tableid 0xe0000000 (0x8cf5b368), Vrfid 0x60000000, Vrid 0x20000000, Flags 0x1019
Vrfname default, Refcount 4129
```

```

56 routes, 0 protected, 0 reresolve, 0 unresolved (0 old, 0 new), 7616 bytes
13 rib, 0 lsd, 0:27 aib, 1 internal, 10 interface, 4 special, 1 default routes
56 load sharing elements, 24304 bytes, 1 references
1 shared load sharing elements, 432 bytes
55 exclusive load sharing elements, 23872 bytes
0 route delete cache elements
13 local route bufs received, 1 remote route bufs received, 0 mix bufs received
13 local routes, 0 remote routes
13 total local route updates processed
0 total remote route updates processed
0 pkts pre-routed to cust card
0 pkts pre-routed to rp 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 LISP EID prefixes, 0 merged, via 0 rlocs
28 next hops
1 incomplete next hop

0 PD backwalks on LDIs with backup path

```

- `show cef { ipv4 address | ipv6 address } detail`

Displays the details of the IPv4 or IPv6 CEF table.

```

Router#show cef 203.0.1.2 detail
203.0.1.2/32, version 0, internal 0x1020001 0x0 (ptr 0x8d7db7f0) [1], 0x0 (0x8daeef0), 0x0
(0x0)
Updated Nov 20 13:33:23.556
local adjacency 203.0.1.2
Prefix Len 32, traffic index 0, Adjacency-prefix, precedence n/a, priority 15
gateway array (0x8d84beb0) reference count 1, flags 0x0, source aib (10), 0 backups
[2 type 3 flags 0x8401 (0x8d99a598) ext 0x0 (0x0)]
LW-LDI[type=3, refc=1, ptr=0x8daeef0, sh-ldi=0x8d99a598]
gateway array update type-time 1 Nov 20 13:33:23.556
LDI Update time Nov 20 13:33:23.556
LW-LDI-TS Nov 20 13:33:23.556
via 203.0.1.2/32, HundredGigE0/0/0/9, 3 dependencies, weight 0, class 0 [flags 0x0]
path-idx 0 NHID 0x0 [0x8cfc81a0 0x0]
next hop 203.0.1.2/32
local adjacency
Load distribution: 0 (refcount 2)

Hash OK Interface Address
0 Y HundredGigE0/0/0/9 203.0.1.2

```

- `show adjacency detail`

Displays detailed adjacency information, including Layer 2 information for each interface. The output of the `show adjacency` command varies by location.

```

Router#show adjacency detail

-----
0/5/CPU0
-----
Interface Address Version Refcount Protocol
Hu0/5/0/12 (interface) 13 1( 0)

```

| | | | |
|------------|--|----|-------|
| | (interface entry) mtu: 1500, flags 1 4 | | |
| Hu0/5/0/30 | (interface) (interface entry) mtu: 1500, flags 1 4 | 31 | 1(0) |
| Hu0/5/0/19 | (interface) (interface entry) mtu: 1500, flags 1 4 | 20 | 1(0) |
| Hu0/5/0/16 | (interface) (interface entry) mtu: 1500, flags 1 4 | 17 | 1(0) |
| Hu0/5/0/23 | (interface) (interface entry) mtu: 1500, flags 1 4 | 24 | 1(0) |
| Hu0/5/0/22 | (interface) (interface entry) mtu: 1500, flags 1 4 | 23 | 1(0) |
| Hu0/5/0/1 | (interface) (interface entry) mtu: 1500, flags 1 4 | 2 | 1(0) |
| Hu0/5/0/6 | (interface) (interface entry) mtu: 1500, flags 1 4 | 7 | 1(0) |
| Hu0/5/0/10 | (interface) (interface entry) mtu: 1500, flags 1 4 | 11 | 1(0) |
| Hu0/5/0/31 | (interface) (interface entry) mtu: 1500, flags 1 4 | 32 | 1(0) |
| Hu0/5/0/28 | (interface) (interface entry) mtu: 1500, flags 1 4 | 29 | 1(0) |
| Hu0/5/0/35 | (interface) (interface entry) mtu: 1500, flags 1 4 | 36 | 1(0) |
| Hu0/5/0/32 | (interface) (interface entry) mtu: 1500, flags 1 4 | 33 | 1(0) |

| | | | |
|------------|--|----|-------------|
| Hu0/5/0/15 | (interface) (interface entry) mtu: 1500, flags 1 4 | 16 | 1 (0) |
| Hu0/5/0/34 | (interface) (interface entry) mtu: 1500, flags 1 4 | 35 | 1 (0) |
| Hu0/5/0/2 | (interface) (interface entry) mtu: 1500, flags 1 4 | 3 | 1 (0) |
| Hu0/5/0/26 | (interface) (interface entry) mtu: 1500, flags 1 4 | 27 | 1 (0) |
| Hu0/0/0/9 | 203.0.1.2 0010940000cea285f0b80248847 mtu: 8986, flags 1 0 | 50 | 2 (0) mpls |
| Hu0/0/0/9 | 203.0.1.2 0010940000cea285f0b80240800 mtu: 8986, flags 1 0 | 49 | 2 (0) ipv4 |
| Hu0/5/0/29 | (interface) (interface entry) mtu: 1500, flags 1 4 | 30 | 1 (0) |
| Hu0/5/0/33 | (interface) (interface entry) mtu: 1500, flags 1 4 | 34 | 1 (0) |
| Hu0/5/0/20 | (interface) (interface entry) mtu: 1500, flags 1 4 | 21 | 1 (0) |
| Hu0/5/0/24 | (interface) (interface entry) mtu: 1500, flags 1 4 | 25 | 1 (0) |
| Hu0/5/0/0 | (interface) (interface entry) mtu: 1500, flags 1 4 | 1 | 1 (0) |
| Hu0/5/0/4 | (interface) (interface entry) mtu: 1500, flags 1 4 | 5 | 1 (0) |
| Hu0/5/0/8 | (interface) (interface entry) mtu: 1500, flags 1 4 | 9 | 1 (0) |

```

Hu0/5/0/3          (interface)          4          1( 0)
                   (interface entry)
                   mtu: 1500, flags 1 4

Hu0/5/0/27         (interface)          28         1( 0)
                   (interface entry)
                   mtu: 1500, flags 1 4

Hu0/5/0/7          (interface)          8          1( 0)
                   (interface entry)
                   mtu: 1500, flags 1 4

Hu0/5/0/14         (interface)          15         1( 0)
                   (interface entry)
                   mtu: 1500, flags 1 4

Hu0/5/0/11         (interface)          12         1( 0)
                   (interface entry)
                   mtu: 1500, flags 1 4

Hu0/5/0/18         (interface)          19         1( 0)
                   (interface entry)

```

Per-Flow Load Balancing

The system inherently supports the 7-tuple hash algorithm. 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.
- Incoming data traffic is evenly distributed over multiple equal-cost connections member links within a bundle interface.
- Layer 2 bundle and Layer 3 (network layer) load balancing decisions are taken on IPv4, IPv6, and MPLS flows. If it is an IPv4 or an IPv6 payload, then a 7-tuple hashing is done. If it is an MPLS payload with three or less labels, then the hardware parses the payload underneath and identifies whether the payload packet has an IPv4 or an IPv6 header. If it is an IPv4 or IPv6 header, then a 4-tuple hashing is performed based on the IP source, IP destination, router ID, and label stack; otherwise, an MPLS label based hashing is performed. In case of MPLS label-based hashing, the top 4 labels are used in hash computation.
- A 7-tuple hash algorithm provides more granular load balancing and 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 7-tuple load-balance hash calculation contains:
 - Source IP address
 - Destination IP address

- IP Protocol type
- Router ID
- Source port
- Destination port
- Input interface

Per-Destination Load Balancing


Per destination load balancing is used for packets that transit over a recursive MPLS path (for example, learned through BGP 3107). Per-destination load balancing means the router distributes the packets based on the destination of the route. Given two paths to the same network, all packets for destination1 on that network go over the first path, all packets for destination2 on that network go over the second path, and so on. This preserves packet order, with potential unequal usage of the links. If one host receives the majority of the traffic all packets use one link, which leaves bandwidth on other links unused. A larger number of destination addresses leads to more equally used links.

Configuring Static Route

Routers forward packets using either route information from route table entries that you manually configure or the route information that is calculated using dynamic routing algorithms. Static routes, which define explicit paths between two routers, cannot be automatically updated; you must manually reconfigure static routes when network changes occur. Static routes use less bandwidth than dynamic routes. Use static routes where network traffic is predictable and where the network design is simple. You should not use static routes in large, constantly changing networks because static routes cannot react to network changes. Most networks use dynamic routes to communicate between routers but might have one or two static routes configured for special cases. Static routes are also useful for specifying a gateway of last resort (a default router to which all unroutable packets are sent).

Configuration Example

Create a static route between Router A and B over a HundredGigE interface. The destination IP address is 203.0.1.2/32 and the next hop address is 1.0.0.2.



```
Router(config)#router static address-family ipv4 unicast
Router(config-static-afi)#203.0.1.2/32 HundredGigE 0/0/0/9 1.0.0.2
Router(config-static-afi)#commit
```

Running Configuration

```
Router#show running-config router static address-family ipv4 unicast
router static
 address-family ipv4 unicast
 203.0.1.2/32 HundredGigE 0/0/0/9 1.0.0.2
!
```

Verification

Verify that the Next Hop Flags fields indicate COMPLETE for accurate functioning of the configuration.

```
Router#show cef 203.0.1.2/32 hardware egress detail location 0/RP0/CPU0
203.0.1.2/32, version 0, internal 0x1020001 0x0 (ptr 0x8d7db7f0) [1], 0x0 (0x8daeedf0), 0x0
(0x0)
Updated Nov 20 13:33:23.557
local adjacency 203.0.1.2
Prefix Len 32, traffic index 0, Adjacency-prefix, precedence n/a, priority 15
  via 203.0.1.2/32, HundredGigE0/0/0/9, 3 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0x8cfc81a0 0x0]
    next hop 1.0.0.2
    local adjacency
```

Associated Commands

- router static
- [show cef](#)

BGP Attributes Download

The BGP Attributes Download feature enables you to display the installed BGP attributes in CEF.

- The **show cef bgp-attribute** command displays the installed BGP attributes in CEF.
- The **show cef bgp-attribute attribute-id** command and the **show cef bgp-attribute local-attribute-id** command are used to view the specific BGP attributes by attribute ID and local attribute ID.

Verification

```
Router# show cef bgp-attribute
Router ID is 216.1.1.1
```

```
IP CEF with switching (Table Version 0) for node0_RP0_CPU0
```

```
Load balancing: L4
Tableid 0xe0000000 (0x8cf5b368), Vrfid 0x60000000, Vrid 0x20000000, Flags 0x1019
Vrfname default, Refcount 4129
56 routes, 0 protected, 0 reresolve, 0 unresolved (0 old, 0 new), 7616 bytes
13 rib, 0 lsd, 0:27 aib, 1 internal, 10 interface, 4 special, 1 default routes
56 load sharing elements, 24304 bytes, 1 references
1 shared load sharing elements, 432 bytes
55 exclusive load sharing elements, 23872 bytes
0 route delete cache elements
13 local route bufs received, 1 remote route bufs received, 0 mix bufs received
13 local routes, 0 remote routes
13 total local route updates processed
0 total remote route updates processed
0 pkts pre-routed to cust card
0 pkts pre-routed to rp 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 LISP EID prefixes, 0 merged, via 0 rlocs
```



```
28 next hops
 1 incomplete next hop

0 PD backwalks on LDIs with backup path
```

VRF: default

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

Associated Commands

- [show cef bgp-attribute](#)

