



Implementing Cisco Express Forwarding

- [Implementing Cisco Express Forwarding, on page 1](#)
- [Prerequisites for Implementing Cisco Express Forwarding, on page 2](#)
- [Verifying CEF, on page 2](#)
- [Configuration Status of Cisco Express Forwarding Hardware Modules, on page 4](#)
- [Configuring Static Route, on page 5](#)
- [BGP Attributes Download, on page 6](#)
- [Proactive Address Resolution Protocol and Neighbor Discovery, on page 6](#)
- [Route Scale Improvements, on page 7](#)
- [Shortened IPv6 Routing Prefixes, on page 8](#)
- [Per-NPU hash rotation, on page 11](#)

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 8000 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

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.

Verifying CEF

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

- `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 102239, internal 0x1000001 0x0 (ptr 0xa932b408) [1], 0x0 (0xaf4a6ad8),
0xa20 (0xc22c6da8)
Updated Jul 3 21:40:17.827
local adjacency 203.1.104.2
Prefix Len 32, traffic index 0, precedence n/a, priority 3
gateway array (0xb9061e70) reference count 1982, flags 0x8068, source lsd (5), 1 backups
[1983 type 4 flags 0x108401 (0x943df068) ext 0x0 (0x0)]
LW-LDI[type=1, refc=1, ptr=0xaf4a6ad8, sh-ldi=0x943df068]
gateway array update type-time 1 Jul 3 20:23:36.957
LDI Update time Jul 3 20:23:36.964
LW-LDI-TS Jul 3 21:40:17.834
via 203.1.104.2/32, Bundle-Ether104, 11 dependencies, weight 0, class 0 [flags 0x0]
path-idx 0 NHID 0x0 [0xa446b0a8 0x0]
next hop 203.1.104.2/32
local adjacency
via 203.1.114.2/32, Bundle-Ether114, 9 dependencies, weight 0, class 0 [flags 0x0]
path-idx 1 NHID 0x0 [0xa446ac18 0x0]
next hop 203.1.114.2/32
local adjacency

Load distribution: 0 1 (refcount 1983)

Hash OK Interface Address

```

```

0      Y      Bundle-Ether104      203.1.104.2
1      Y      Bundle-Ether114     203.1.114.2

```

Configuration Status of Cisco Express Forwarding Hardware Modules

Table 1: Feature History Table

Feature name	Release Information	Feature Description
Configuration Status of Cisco Express Forwarding (CEF) Hardware Modules	Release 24.4.1	Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*) *This feature is supported on: <ul style="list-style-type: none"> • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM
Configuration Status of Cisco Express Forwarding (CEF) Hardware Modules	Release 24.1.1	Introduced in this release on: Cisco 8011-4G24Y4H-I router.
Configuration Status of Cisco Express Forwarding (CEF) Hardware Modules	Release 7.3.1	This feature enables you to view pending actions, such as a reload or a commit action, which is applicable to CEF hardware-modules. The show hw-module profile cef command is introduced for this feature.

To understand the configuration status of different CEF hardware modules, use the **show hw-module profile cef** command. Use the command in the XR EXEC mode to display the following information:

```

Router# show hw-module profile cef
Tue Oct 6 00:34:47.735 UTC

```

```

-----
Knob                               Status      Applied    Action
-----
BGPLU                              Configured  No         Reload
Dark Bandwidth                     Unconfigured Yes        None
MPLS Per Path Stats                Unconfigured Yes        None

```

Tunnel TTL Decrement	Configured	Yes	None
High-Scale No-LDP-Over-TE	Unconfigured	Yes	None

You can get the following information from the output of the **show hw-module profile cef** command:

- The **Status** column indicates whether you have configured the corresponding hardware-module for CEF.
- The **Applied** column indicates whether you have applied the command on the router.
- The **Action** column indicates whether you must reload router to apply the configuration.

Restrictions

The Yang data model for the **show hw-module profile cef** command is not available.

Associated Commands

- [show hw-module profile cef](#)

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 203.0.113.0/24 and the next hop address is 192.0.2.1.



```
Router(config)#router static address-family ipv4 unicast
Router(config-static-afi)#203.0.113.0/24 HundredGigE0/0/0/0 192.0.2.1
Router(config-static-afi)#commit
```

Running Configuration

```
Router#show running-config router static address-family ipv4 unicast
router static
 address-family ipv4 unicast
 203.0.113.0/24 HundredGigE0/0/0/0 192.0.2.1
 !
 !
```

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
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
```

Proactive Address Resolution Protocol and Neighbor Discovery

When CEF installs a route for which there is no layer 2 adjacency information, CEF creates an incomplete layer 3 next-hop and programs it on the hardware. Because of this incomplete programming, the first packet will be forwarded to the software forwarding path. The software forwarding in turn strips off the layer 2 header from the packet and forwards it to ARP (Address Resolution Protocol) or ND (Neighbor Discovery) in order to resolve the layer 2 adjacency information. In such a packet, if there is feature specific information present in the layer 2 header, the software forwarding path fails to strip off the layer 2 header completely and thus ARP or ND is unable to resolve the missing layer 2 adjacency information and thereby this results in traffic being dropped.

Proactive ARP and ND feature solves the above problem by ensuring that CEF proactively triggers ARP or ND in order to resolve the missing layer 2 adjacency information, retrying every 15 seconds until the next-hop information is resolved. Thus, when you configure a static route which has an incomplete next-hop information, this feature automatically triggers ARP or ND resolution.

Configuration

```
/* Enter the configuration mode and configure Proactive ARP/ND */
Router# configure
Router(config)# cef proactive-arp-nd enable
Router(config)# commit
```

Running Configuration

```
Show running-config
cef proactive-arp-nd enable
end
```

Route Scale Improvements

Table 2: Feature History Table

Feature Name	Release Information	Description
Route Scale Improvements	Release 7.9.1	<p>This feature enables you to increase the number of Forwarding Information Base (FIB) entries supported for IPv4 traffic from 2 million to 3 million and IPv6 traffic from 0.5 million to 1 million.</p> <p>The increased FIB entries allow the router to route more traffic streams. It also helps the router to achieve a faster switch or process-switch forwarding scenario by eliminating the frequent need for route cache maintenance due to fewer route entries in the FIB database.</p> <p>This feature introduces the hw-module profile route scale command.</p>

The Forwarding Information Base (FIB) process maintains the forwarding database for IPv4 and IPv6 unicast in the route processor (RP) and line card (LC). Each routing protocol selects its set of suitable routes and installs those routes and their attributes in the Routing Information Base (RIB). RIB stores these routes and uses it to select an appropriate route for each routing protocol. The FIB on each RP or LC downloads those routes to the line cards and uses them for forwarding packets. The route scale is the capacity of the FIB database to store route entries for IPv4 and IPv6 traffic types.

With this feature, you can increase the routing scale for IPv4 and IPv6 traffic types. For IPv4, you can increase the FIB resources from 2 million to 3 million entries. In IPv6, you can increase the FIB resources from 0.5 million to 1 million entries. The router supports this scale improvement by reassigning the available resources within itself. Increasing the routing scale restricts resources for packet classification features such as Security ACL, QoS ACL, BGP Flowspec, and LPTS.

Restrictions

Enabling the route scale improvements results in resource crunch for packet classification features such as Security ACL, QoS ACL, BGP Flowspec, and LPTS.



Note Enabling route scale improvements for IPv4 and IPv6 traffic types, reduces the TCAM space for other packet classification features on the router by 512 and 256 entries respectively.

Configuration

The following sections details the configuration to increase the route scale for IPv4 traffic type to 3 million and IPv6 traffic type to 1 million entries:

```
Router # config
Router(config)# hw-module profile route scale lpm tcam-banks
Router(config)# commit
Router(config)# reload location all
```



Note For restoring the route scale for IPv4 traffic type to 2 million and IPv6 traffic type to 0.5 million entries, use the **no** form of the **hw-module profile route scale lpm tcam-banks** command and execute the **reload location all** command.

Running Configuration

```
Router# show running-config
Building configuration...
hw-module profile route scale lpm tcam-banks
```

Shortened IPv6 Routing Prefixes

Table 3: Feature History Table

Feature Name	Release Information	Description
Shortened IPv6 Routing Prefixes	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100])(select variants only*); Modular Systems (8800 [LC ASIC: P100])(select variants only*)</p> <p>*This feature is supported on:</p> <ul style="list-style-type: none"> • 8212-48FH-M • 8711-32FH-M • 88-LC1-36EH • 88-LC1-12TH24FH-E • 88-LC1-52Y8H-EM

Feature Name	Release Information	Description
Shortened IPv6 Routing Prefixes	Release 24.1.1	<p>We've improved the memory utilization of the Longest Prefix Match (LPM) table, allowing you to accommodate more IPv6 routing prefixes with wider IPv6 prefix lengths by storing them as shorter routing prefixes. This approach conserves router resources by reducing the number of entries in the routing table and is particularly helpful in mitigating out-of-resource (OOR) situations.</p> <p>This feature is supported only on Cisco Silicon One Q200-based routers and line cards.</p> <p>This feature modifies the hw-module profile route scale command with the new lpm wide-entries shortened keyword.</p>

Overview

An IPv6 address is 128-bits long and is composed of three main parts – the routing prefix, subnet ID, and interface ID. The bits of the routing prefix of an IPv6 address are stored in the Longest Prefix Match (LPM) table. The router uses the longest prefix match entry to determine the egress interface and the address of the next destination device for a packet.



Note The **hw-module profile route scale lpm wide-entries shortened** command isn't enabled by default, and we recommend using it judiciously to accommodate higher number of wide-entry IPv6 prefixes.

From Release 24.1.1 onwards, when you configure the **hw-module profile route scale lpm wide-entries shortened** command, if the size of the subnet prefix is more than 64-bits, those wide routing prefixes are stored in the LPM table in the form of shorter routing prefixes. This allows you to accommodate more wide-route prefixes in the LPM table, hence allowing you to mitigate Out Of Resource (OOR) situation.

Configuration Guidelines and Limitations

Following are the configuration guidelines and restrictions for the shortened IPv6 routing prefix feature:

- Router reload is required after you enable this feature.
- This feature is supported only on the Cisco Silicon One Q200-based routers and line cards.
- [IPv6 Prefix Scale Expansion](#) isn't supported.
- Unicast Reverse Packet Forwarding (uRPF) isn't supported.

Configuration

The following section details the configuration to enable shortening for the wide routing prefixes:

```
Router# configure
Router(config)# hw-module profile route scale lpm wide-entries shortened
Router(config)# commit
Router# reload location all
```



Note To disable this feature, use the **no** form of the **hw-module profile route scale lpm wide-entries shortened** command and execute the **reload location all** command.

Running Configuration

```
Router# show running-config
Building configuration...
hw-module profile route scale lpm wide-entries shortened
```

Verification

Use the **show hw-module profile route-scale** command to verify the **Wide-Entries- Shortened** configuration.

```
Router#show hw-module profile route-scale
Thu Feb  8 17:06:09.652 UTC
-----
Knob                               Status           Applied   Action
-----
IPv6-Pfx-Expansion                 Unconfigured    N/A       None
LPM-Tcam-Scale                     Unconfigured    N/A       None
Wide-Entries-Shortened           Configured    Yes     None
```

Per-NPU hash rotation

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
Per-NPU hash rotation	Release 24.4.1	<p>Introduced in this release on: Fixed Systems (8200 [ASIC: Q200, P100], 8700 [ASIC: P100, K100]); Centralized Systems (8600 [ASIC:Q200]); Modular Systems (8800 [LC ASIC: Q100, Q200, P100])</p> <p>You can now configure the hash rotation value for each Network Processing Unit (NPU) to improve traffic load balancing and minimize traffic polarization. Alternatively, the value can be automatically calculated, eliminating the need for manual configuration.</p>

The per-NPU hash rotation is a technique that is used to mitigate the challenge of the traffic polarization in a highly symmetric network on Silicon One devices by adjusting the hash rotate value which alters the hashing function using the **cef load-balancing algorithm adjust** command.

The **cef load-balancing algorithm adjust** command can be configured either in global-level or per-NPU-level.

- Global-level Configuration: To configure per-ASIC settings, use the **cef load-balancing algorithm adjust <value>** command.
- Per-NPU-level Configuration: The router can automatically calculate and assign different hash rotation values randomly to each NPU. This approach provides varied hash values within the router and the network algorithm to mitigate traffic polarization. These values can be overridden using the CLI command **cef load-balancing algorithm adjust <value> instance <npu_id> location <LC>** for individual NPUs, or **cef load-balancing algorithm adjust <value>** at the global level, giving administrators control over load-balancing behavior.

Key attributes

Hash function: Hash function processes input to produce a fixed-size string of bytes, known as hash code. It plays a crucial role in determining the path or interface that a packet takes, balancing the traffic load across multiple paths.

Traffic polarization: When a network's hash function consistently maps certain types of traffic to the same path, leading to uneven distribution and potential congestion on that path while other paths remain underutilized.

Hash rotation: Hash rotation involves modifying the hash function or its parameters to change the distribution of hash values. This determines the path or interface that a packet should take based on its header information. By rotating hash values either at each router level in the network, or per-NPU level within each router in a nonuniform or random manner, network devices can achieve a more balanced traffic distribution across

multiple paths or interfaces. This leads to improved load balancing, reduced traffic polarization, effective link utilization, and decreased congestion.

Restrictions

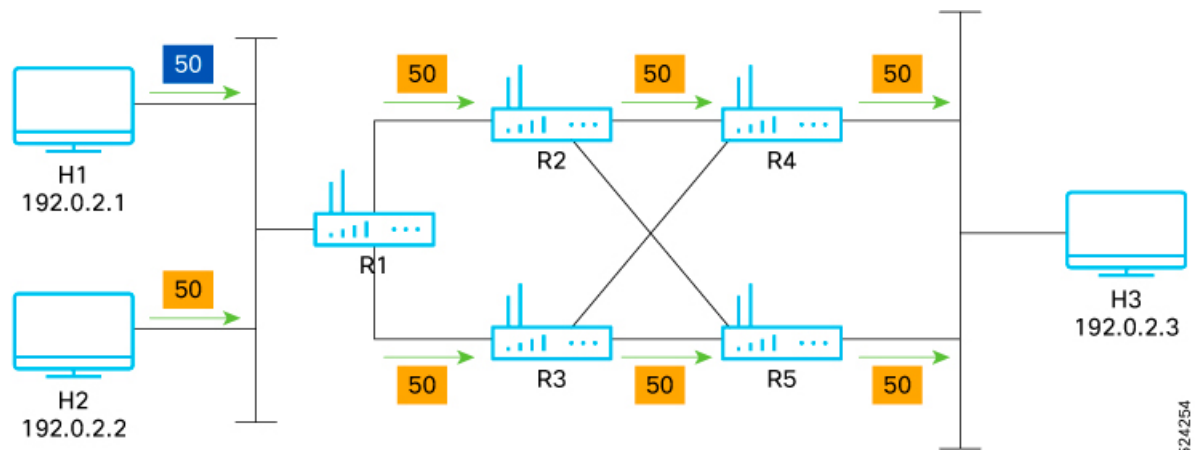
Use this feature only when you notice traffic polarization and underutilized bandwidth.

Traffic distribution with hash rotation

In a network with multiple paths between two nodes, a hash function determines which path each packet takes. If the hash function consistently maps certain types of traffic to the same path, that path can become congested. By configuring hash rotation, the network device adjusts the hash function to distribute the traffic more evenly across all the available paths.

Figure 1: Network topology for optimized hash rotation and load balancing

This figure shows the network topology for hash rotation, illustrating how routers R1, R2, and R3 distribute packets from nodes H1 and H2 to balance load and optimize path utilization.



Packets from hosts H1 and H2 are load balanced by router R1. When the packets reach routers R2 and R3:

- Router R2 processes packets from host H1 using the same hashing algorithm as router R1. Because all packets are destined for host H3, they are routed through a single interface, leaving the link between routers R2 and R5 unused.
- Router R3 receives the packets from host H2 and runs the same hashing algorithm as router R1. Because all packets are destined for host H3, they are routed through a single interface, leaving the link between routers R3 and R4 unused.

This scenario exemplifies Cisco Express Forwarding (CEF) polarization, where using the same hashing algorithm and inputs across all routers results in a single path being used for specific packet flow.

Configure per-NPU hash rotation

You can adjust the hash rotation value to modify the hashing function in highly symmetrical networks, effectively addressing traffic polarization and improving load balancing.

Before you begin

Enable the **hw-module profile pbr vrf-redirect** command and reload the router.

Procedure

You can configure hash rotation in one of these ways. This defines the priority order in which the hash rotation value is programmed in the hardware.

Example:

a. Per-NPU hash rotate value configuration

This allows a unique hash rotate value to be programmed in hardware per-NPU.

```
Router# configure
Router(config)# cef load-balancing algorithm adjust 7 instance 0 location 0/0/CPU0
Router# end
```

Hash rotate count ranges 1—35. The default value is 1 when the **hw-module profile pbr vrf-redirect** command is enabled.

b. Global (system-level) configuration

This sets a global hash rotate value for all NPUs.

```
Router# configure
Router(config)# cef load-balancing algorithm adjust 2
Router# end
```

c. Automatic hash rotate calculation

This method automatically calculates the hash rotation value using the router ID, line card ID, and NPU ID when the **hw-module profile pbr vrf-redirect** command is enabled, requiring no additional configuration.

If	then use the
you want to know the values that are programmed in the hardware for all NPUs of a given line card	<code>show controllers npu debugshell <npu-id> 'script device_hash_rotate_info get_val_all_npu' location <LC_node></code> command.
you want to know the values that are programmed in the hardware for a specific NPU of a given line card	<code>show controllers npu debugshell <npu_id> 'script device_hash_rotate_info get_val' location <LC_node></code> command.

