

Understand and Troubleshoot CEF on Cisco IOS XE Routers

Contents

[Introduction](#)

[CEF Behavior on the Cisco IOS XE Platform](#)

[Check CEF Adjacency](#)

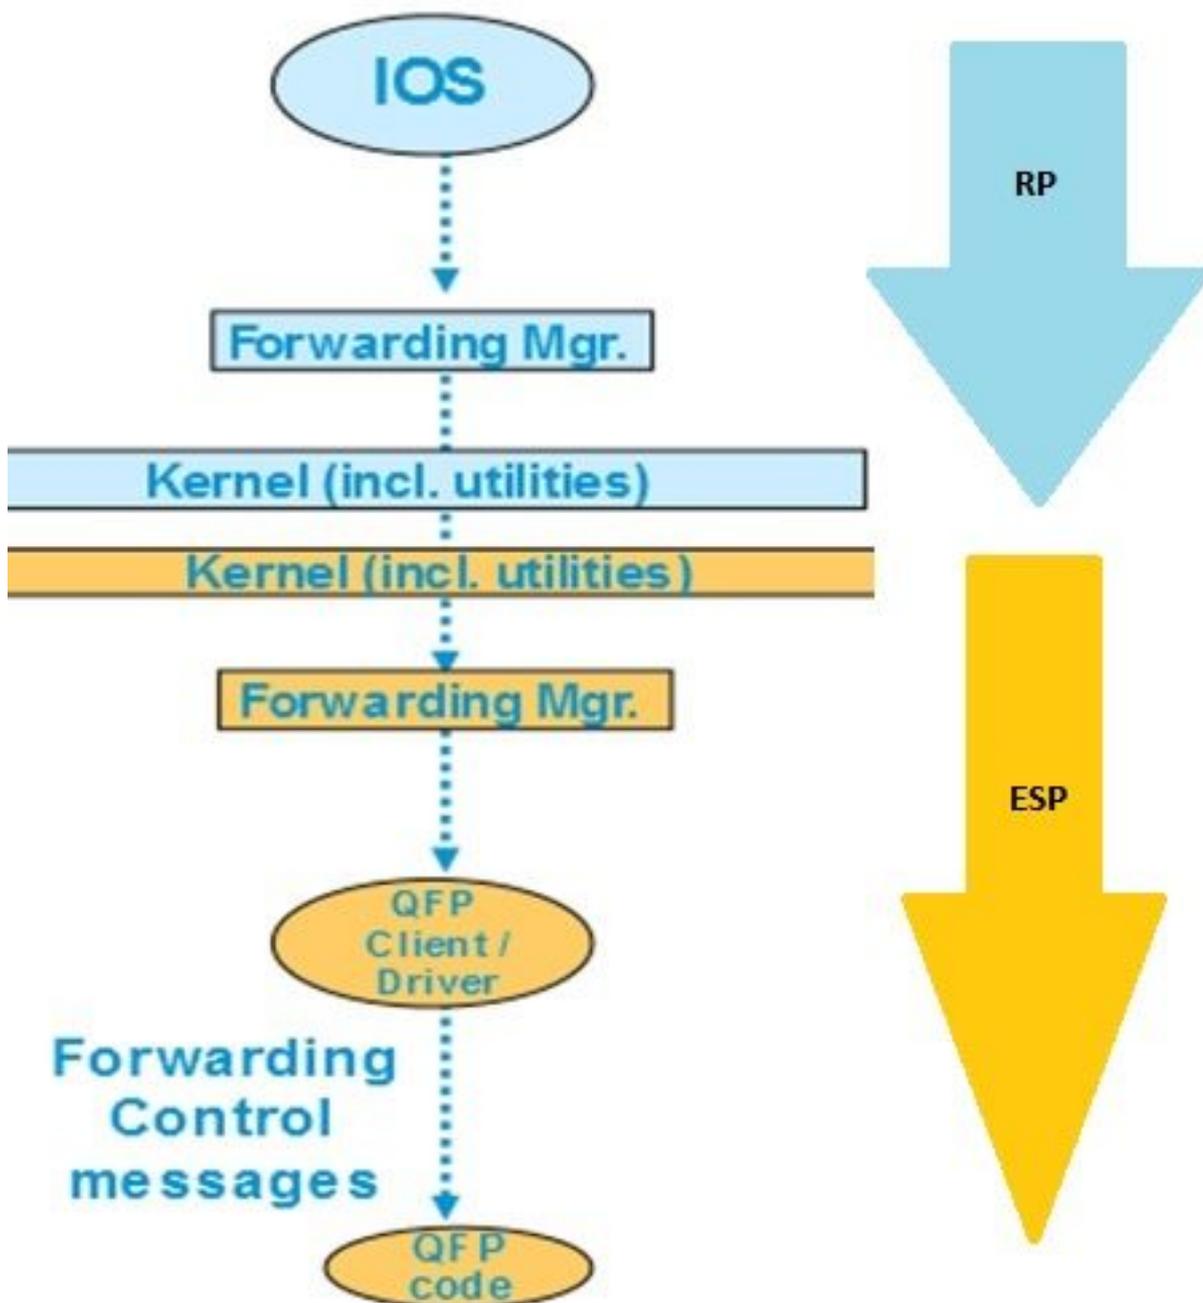
[Common Phenomenon Observed](#)

[Conclusion](#)

Introduction

This document describes the Cisco Express Forwarding (CEF) feature on Cisco IOS[®] XE based devices. Unlike other Cisco routers, Cisco IOS XE-based routers are modular in nature not only in terms of hardware, but also in software. Due to this nature, the behavior of most of the features and protocols are also a little different. You will also see how CEF tables are maintained on Cisco IOS XE-based devices and how big Border Gateway Protocol (BGP) tables are managed in terms of CEF updates on Cisco IOS XE platforms.

CEF Behavior on the Cisco IOS XE Platform



CEF table

update inside XE platform

On Cisco IOS XE devices such as the ASR1000, the control plane is separate to the forwarding plane. Whenever any update needs to be passed from the control plane to the data plane, it has to go through the data flow shown in the flow diagram. For example, in case of CEF whenever any prefix is learned on the control plane, this update passes from the control plane (IOSd) to the forwarding manager of the control plane (FMAN-RP). The forwarding manager on the control plane uses kernel utilities like Ismpi, Hyper-transport (HT) links, and so on in order to pass the update to forwarding plane (ESP's) forwarding manager (FMAN-FP). The forwarding manager sends the update to the Quantum Flow Processor (QFP) which programs QFP microcode in order to finally program the QFP subsystem which does the actual forwarding of packets in Cisco Aggregation Services Router (ASR) devices.

There are various commands you can use to check the CEF update on each of these software modules. This is the step-by-step process for that.

In order to check CEF on the control plane:

```
Router#show ip cef Prefix Next Hop Interface
0.0.0.0/0 no route 0.0.0.0/8 drop 0.0.0.0/32
receive 1.1.1.1/32 10.10.10.1 GigabitEthernet0/0/0
2.2.2.2/32 receive Loopback1 10.10.10.0/24
attached GigabitEthernet0/0/0 10.10.10.0/32 receive
GigabitEthernet0/0/0 Router#show platform software ip rp active cef summary Forwarding Table
Summary Name VRF id Table id Protocol Prefixes State -----
-----
Default 0 0 IPv4 20 OM handle: 0x404a4df8
Router#show platform software ip rp active cef detail Forwarding Table 0.0.0.0/0 ->
OBJ_ADJ_NOROUTE (0), urpf: 5 Prefix Flags: Default, Default route handler OM handle: 0x404a91e8
0.0.0.0/8 -> OBJ_ADJ_DROP (0), urpf: 13 Prefix Flags: unknown OM handle: 0x404bd5e8 0.0.0.0/32 -
-> OBJ_ADJ_RECEIVE (0), urpf: 12 Prefix Flags: Receive OM handle: 0x404bd298 1.1.1.1/32 ->
OBJ_ADJ_AGENCY (16), urpf: 20 Prefix Flags: unknown OM handle: 0x404fec70
```

In order to check the CEF details in the forwarding plane (ESP):

```
Router#show platform software ip fp active cef detail Forwarding Table 0.0.0.0/0 ->
OBJ_ADJ_NOROUTE (0), urpf: 5 Prefix Flags: Default, Default route handler aom id: 73, HW handle:
0x4310df8 (created) 0.0.0.0/8 -> OBJ_ADJ_DROP (0), urpf: 13 Prefix Flags: unknown aom id: 90, HW
handle: 0x4362cd8 (created) 0.0.0.0/32 -> OBJ_ADJ_RECEIVE (0), urpf: 12 Prefix Flags: Receive
aom id: 86, HW handle: 0x4333568 (created) 127.0.0.0/8 -> OBJ_ADJ_DROP (0), urpf: 13 Prefix
Flags: unknown aom id: 91, HW handle: 0x4387048 (created) 224.0.0.0/4 -> OBJ_ADJ_DROP (0), urpf:
13 Prefix Flags: unknown aom id: 92, HW handle: 0x43870d8 (created) Router#show platform
software ip fp active cef summary Forwarding Table Summary Name VRF id Table id
Protocol Prefixes State -----
-----
Default 0 0 IPv4
20 hw: 0x43010a8 (created)
```

These commands can also be used when you face CEF issues on the device. For example, although the routes are learned, the prefixes are not reachable. You can dig through all the modules to see if all the CEF tables are updated properly or not.

Check CEF Adjacency

In a similar way, you can further check the CEF adjacency table for all the Layer 2 information about the adjacent prefixes.

In order to check CEF adjacency on the control plane:

```
Router#show adjacency gigabitEthernet 0/0/0 detail Protocol Interface Address
IP GigabitEthernet0/0/0 10.10.10.1(11) 72772
packets, 4622727 bytes epoch 0
sourced in sev-epoch 0
Encap length 14 0062EC6B89000062EC6BEC000800
L2 destination address byte offset 0
L2 destination address byte length 6
Link-type after encap: ip
ARP Router#show platform software adjacency rp active Number of adjacency objects: 4 Adjacency
id: 0x10 (16) Interface: GigabitEthernet0/0/0, IF index: 8, Link Type: MCP_LINK_IP Encap:
0:62:ec:6b:89:0:0:62:ec:6b:ec:0:8:0 Encap Length: 14, Encap Type: MCP_ET_ARPA, MTU: 1500
Flags: no-l3-inject Incomplete behavior type: None Fixup: unknown Fixup_Flags_2: unknown
Nexthop addr: 10.10.10.1 IP FRR MCP_ADJ_IPFRR_NONE 0 OM handle: 0x404ea1d8
```

You need to note the adjacency ID in order to check the details about this particular adjacency in the forwarding plane. In this case, the **Adjacency ID is 16**.

In order to check CEF adjacency on the forwarding plane:

```
Router#show platform software adjacency fp active index 16 Number of adjacency objects: 4
Adjacency id: 0x10 (16) Interface: GigabitEthernet0/0/0, IF index: 8, Link Type: MCP_LINK_IP
Encap: 0:62:ec:6b:89:0:0:62:ec:6b:ec:0:8:0 Encap Length: 14, Encap Type: MCP_ET_ARPA, MTU:
1500 Flags: no-l3-inject Incomplete behavior type: None Fixup: unknown Fixup_Flags_2:
unknown Nexthop addr: 10.10.10.1 IP FRR MCP_ADJ_IPFRR_NONE 0 aom id: 114, HW handle:
0x43ae148 (created)
```

Here, you see that the CEF adjacency information is populated in the Forwarding manager (FMAN) on FP. FMAN FP sends this information to the QFP client driver which programs the QFP forwarding table which will be used for forwarding eventually. From the previous command, copy the hardware handle in order to check the forwarding information on QFP.

```
Router#show pla hard qfp act feature cef-mps adjacency handle 0x43ae148 Adj Type: : IPV4
Adjacency Encap Len: : 14 L3 MTU: : 1500 Adj Flags: : 0 Fixup Flags: : 0 Output UIDB: :
Interface Name: GigabitEthernet0/0/0 Encap: : 00 62 ec 6b 89 00 00 62 ec 6b ec 00 08 00 Next Hop
Address: : 10.10.10.1 Lisp Fixup HW Ptr: : 0x767b28f0 Next HW OCE Ptr: : 00000000 CM HW
Ptr: : 946947588 Fixup_Falgs_2: : 0
```

Here, you know that all the adjacency tables are updated properly and the router is forwarding ready. However, the whole process of isolating takes lots of commands and requires knowledge of the modular architecture at a certain level. Hence, in order to simplify this, there was a command introduced recently which gives consolidated information from all the modules.

Note: For the devices with a long routing table, this command might take several minutes to run.

The command is **show ip cef platform detail**.

Common Phenomenon Observed

For all the Cisco IOX XE modular devices in the situations where a huge number of prefixes are learned on the router, normally it takes some time to program all the prefixes in all the forwarding modules. This can be seen very frequently on the routers which are sitting at provider edge learning full BGP routing table from ISP.

In the Technical Assistance Center, there were few cases received where it was seen that after the BGP session comes up and even the BGP route is updated in routing table, the prefixes are not reachable for a while. Normally, it takes 20-30 seconds and it depends on the router platform to ping those prefixes. For example, here is a test scenario:



ASR1002-HX



Pagent running on Cisco 3900

Pagent is a traffic generator tool which is used to push one million BGP routes to the ASR1002HX router.

Here you see that, even if the BGP routes are learned on the device and the control plane CEF table is updated, the internal network is unable to ping the learned prefixes for few more seconds.

On the basis of the CEF discussion, it is clear that you need to have CEF entries updated on each software module. You can see one consequence of this behavior in this particular scenario where the prefixes are not reachable due to the fact that it was not updated in the ESP forwarding table. Here are a few outputs from the ASR1002HX for reference.

BGP tables are updated with all one million routes.

```
Router#show ip bgp summary BGP router identifier 1.1.1.1, local AS number 100 BGP table version
is 1, main routing table version 1 1000002 network entries using 248000496 bytes of memory
1000002 path entries using 128000256 bytes of memory 100002/0 BGP path/bestpath attribute
entries using 26400528 bytes of memory 100000 BGP AS-PATH entries using 5402100 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory 0 BGP filter-list cache entries using 0
bytes of memory BGP using 407803380 total bytes of memory BGP activity 8355774/7355772 prefixes,
9438985/8438983 paths, scan interval 60 secs Neighbor      V          AS  MsgRcvd MsgSent
TblVer  InQ  OutQ          Up/Down          State/PfxRcd
10.10.10.2      4      100      5      2      1      0      0 00:00:58
                1 20.20.20.2      4      100 100002      3
1      0      0 00:01:02      1000000
```

Although, the BGP table has one million prefixes, the forwarding manager CEF table had only **48613** prefixes learned yet.

If you wait for 20-30 seconds, you see the fully updated FP CEF table with one million prefixes.

```
Router#show platform software ip fp active cef summary
Forwarding Table Summary
Name          VRF id  Table id  Protocol  Prefixes  State
-----
Default       0       0         IPv4      48613     hw: 0x2edce98 (created)
```

Conclusion

When you deal with Cisco IOS XE based modular architecture devices for forwarding related issues, you must verify the forwarding table related information from all the software modules. The BGP scenario explained can be considered as expected behavior with this platform as the device takes a few seconds to update the prefixes in all the software modules.