Cisco Express Forwarding

Understanding and troubleshooting CEF in Cisco routers and switches

Nakia Stringfield, CCIE® No. 13451
Russ White, CCIE No. 2635
Stacia McKee

ciscopress.com

Cisco Express Forwarding

Nakia Stringfield, CCIE No. 13451
Russ White, CCIE No. 2635
Stacia McKee

Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

- **224.0.0.0/24 receive**—The range of multicast IP addresses, 224.0.0.0 through 224.0.0.255, is reserved for the use of routing and discovery protocols. For example, Enhanced Interior Gateway Routing Protocol (EIGRP) transmits hello packets to multicast address 224.0.0.10 to maintain neighbor relationships. This FIB entry allows the router to receive and process packets destined to this IP address range.
- **255.255.255.255/32 receive**—This entry allows the router to receive packets transmitted to the link local broadcast address.

Interface-Specific FIB Entries

The three types of interfaces of primary concern when looking at a normal FIB table are as follows:

- Multiaccess interfaces (including broadcast, such as Ethernet)
- Point-to-point interfaces (including most WAN interfaces, such as Packet over SONET [PoS], Frame Relay, and High-Level Data Link Control [HDLC] links)
- Links that are configured to use 31-bit subnet masks

The next sections describe each of these FIB entries.

FIB Entries Built for a Multiaccess Network Interface

A set of FIB entries is built for a multiaccess (or broadcast) interface by default. Example 2-11 shows the set of FIB entries built for an Ethernet interface.

Example 2-11 *FIB Entries Built for a Multiaccess Network Interface*

```
Router#show ip cef
Prefix              Next Hop            Interface
172.16.12.0/24      attached            Ethernet3/0
172.16.12.0/32      receive
172.16.12.3/32      receive
172.16.12.255/32    receive
```

The following list describes the FIB entries associated with a multiaccess network interface:

- **172.16.12.0/24**—This connected FIB entry is built for the network. The subnet prefix points to a glean adjacency. A glean adjacency exists when a specific next hop is directly connected, but no Layer 2 header rewrite string is associated with the entry. Glean adjacencies are covered in the next section.
- **172.16.12.0/32**—This receive FIB entry is built for the all-0s host address on the network segment.

- **172.16.12.3/32**—This receive FIB entry is built for the local interface address so that the router can receive and process packets sent to this IP address.
- **172.16.12.255/32**—This receive FIB entry is built for the all-1s host address on the network segment.

FIB Entries Built on a Point-to-Point Network Interface

A set of FIB table entries are built for point-to-point interfaces. The primary difference between these entries and FIB entries built for a multiaccess interface is that the 192.168.6.0/24 attached FIB entry points to a point-to-point adjacency for the interface, as shown in Example 2-12.

Example 2-12 *FIB Entries Built on a Point-to-Point Network Interface*

```
Router#show ip cef
Prefix              Next Hop            Interface
192.168.6.0/24      attached            Serial5/3
192.168.6.0/32      receive
192.168.6.3/32      receive
192.168.6.255/32    receive
```

FIB Entries Built on a 31-Bit Prefix Network Interface

RFC 3021 permits the use of a 31-bit prefix length on point-to-point links to conserve IP address space. Some FIB entries are not built on these links because no broadcast addresses are on a link with only two hosts. As shown in Example 2-13, an interface configured with a 31-bit mask can only have the connected FIB entry for the network and the receive FIB entry for the locally assigned IP address.

Example 2-13 *FIB Entries Built on a 31-Bit Prefix Network Interface*

```
Router#show ip cef
Prefix              Next Hop            Interface
10.0.6.2/31         attached            Serial5/3
10.0.6.3/32         receive
```

Special Adjacencies

In addition to entries in the FIB, the adjacency table contains special types of adjacencies, including the following:

- Auto
- Punt
- Glean
- Drop

- Discard
- Null
- No route
- Cached and uncached
- Unresolved

The following sections describe each type of adjacency in detail.

Auto Adjacencies

Auto adjacencies are the most common type of adjacencies. Example 2-14 illustrates an auto adjacency installed for interface Serial2/0.

Example 2-14 *Auto Adjacencies*

```
interface Serial2/0
 ip address 10.1.1.1 255.255.255.252

Router#show adjacency internal
Protocol Interface                 Address
IP       Serial2/0                 point2point(5)
                                   0 packets, 0 bytes
                                   0F000800
                                   CEF   expires: 00:02:39
                                         refresh: 00:00:39
                                   Epoch: 0
                                   Fast adjacency disabled
                                   IP redirect enabled
                                   IP mtu 1500 (0x0)
                                   Fixup disabled
                                   Adjacency pointer 0x1BC6D48, refCount 5
                                   Connection Id 0x000000
                                   Bucket 10
```

The protocol type (in this case IP), the Layer 2 header rewrite string, the source of the adjacency entry, and the outbound interface are all included in an adjacency entry. The next hop can be an actual address or it can be marked as point-to-point**,** which indicates that the next hop is reachable over a point-to-point interface.

Punt Adjacency

Punt adjacencies are used when packets to a destination cannot be CEF switched. When a feature is not supported in the CEF switching path, the punt adjacency allows a packet to be switched using the next slower switching mechanism configured on the router. Packets that are punted from the CEF path are sent up to the fast-switching or process-switching path to be handled. All FIB entries that point to a punt adjacency can be viewed using the

Cisco IOS command **show ip cef adjacency punt**. Example 2-15 shows a punt adjacency installed for a multicast address on a router enabled for IP Multicast Routing. IP multicast packets are not supported in the CEF switched path; they are fast- or hardware-switched through a router, depending on the platform type.

Example 2-15 *Punt Adjacency*

```
Router#show ip cef 239.1.1.1
224.0.0.0/4, version 17, epoch 0
0 packets, 0 bytes
  via 0.0.0.0, 0 dependencies
    next hop 0.0.0.0
    valid punt adjacency
```

Glean Adjacency

When a router is directly connected to a segment shared by multiple hosts, the FIB table on the router maintains a prefix for the subnet, rather than an individual prefix for each host. The subnet prefix points to a glean adjacency. When packets need to be forwarded to a specific host, the adjacency database is gleaned for the specific prefix. The FIB entry pointing to a glean adjacency indicates that any addresses within this range should be directly connected to the router, but no Layer 2 header rewrite string is associated with the specific host.

All FIB entries that point to a glean adjacency can be seen using the Cisco IOS command **show ip cef adjacency glean**. The following list and Figure 2-10 explain a simple glean operation:

1. The first packet arrives at Router B from 192.168.1.1 through Router A. Router B performs a FIB lookup for the destination IP address, follows the FIB table down to the longest match, and finds the 172.30.21.0/24 entry. This entry points to a glean adjacency. A glean adjacency causes CEF to punt the packet up to the process level, which in turn hands the packet off to ARP to resolve the Layer 2 header rewrite information for the destination IP address.
2. Router B sends an ARP request for this host destination. The host, 172.30.21.8, responds to the ARP request with its MAC address embedded.
3. Router B receives the ARP reply and builds a new ARP entry.
4. Upon creation of the adjacency, a new FIB ADJFIB entry is created for this host, 172.30.21.8/32, on Router B. Now packets switched through Router B destined for this host IP address can follow the FIB to this new host entry and then to the adjacency table entry, which contains the MAC layer header rewrite string for 172.30.21.8.

Figure 2-10 *Glean Adjacency*

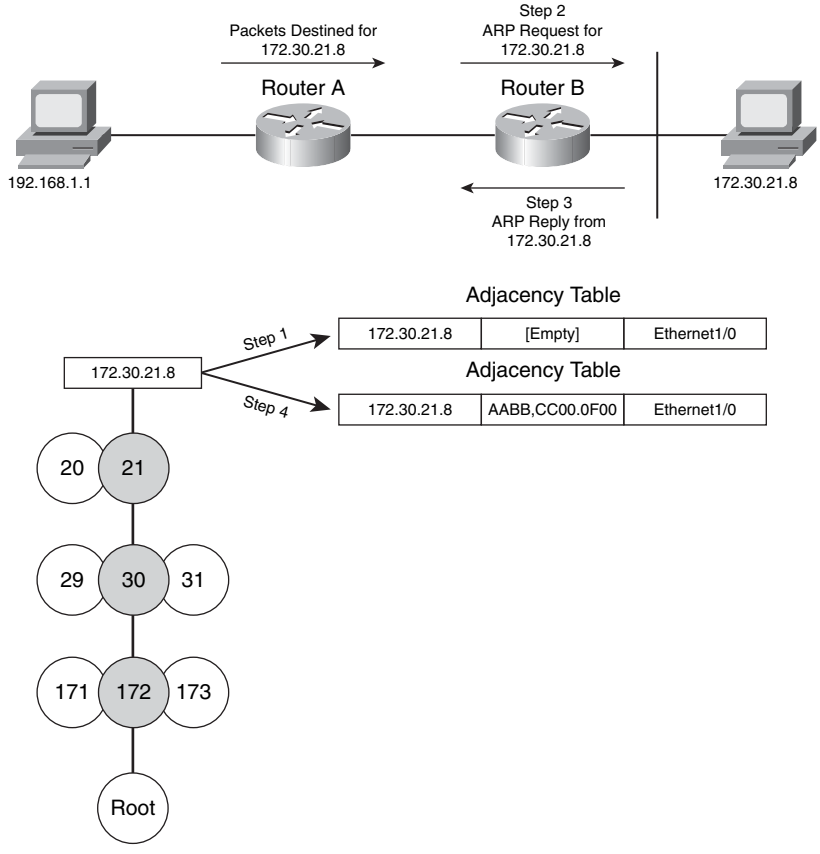

Drop Adjacency

The FIB points to a drop adjacency when packets switched to this destination should be dropped upon being received. Drop adjacencies obey any configured Internet Control Message Protocol (ICMP) requirements. Example 2-16 shows a drop adjacency installed for multicast range 224.0.0.0/4 on a router that does not have IP Multicast Routing enabled.

Example 2-16 *Drop Adjacency*

```
Router#show ip cef 224.0.1.1
224.0.0.0/4, version 19, epoch 0
0 packets, 0 bytes
  via 0.0.0.0, 0 dependencies
    next hop 0.0.0.0
    valid drop adjacency
```

Discard Adjacency

Discard adjacencies are used for addresses that are part of a loopback interface's subnet, but are not actually configured on the loopback interface. In Example 2-17, IP addresses 192.168.1.2 through 192.168.1.254 would all point to a discard adjacency, and packets received by the router, destined to those addresses, would be silently discarded. Packets dropped by a discard adjacency do not generate ICMP messages by default.

Example 2-17 *Discard Adjacency*

```
interface loopback0
 ip address 192.168.1.1 255.255.255.0

Router#show ip cef 192.168.1.2
192.168.1.0/24, version 18, epoch 0, attached, connected
0 packets, 0 bytes
  via Loopback0, 0 dependencies
    valid discard adjacency
```

All FIB entries pointing to a discard adjacency can be viewed using the Cisco IOS command **show ip cef adjacency discard**, as shown in Example 2-18**.**

Example 2-18 **show ip cef adjacency discard** *Command*

```
Router#show ip cef adjacency discard
Prefix              Next Hop            Interface
192.168.1.0/24      attached            Loopback0
```

Null Adjacency

CEF uses null adjacencies to indicate packets that are forwarded to the Null0 interface on a router. Destinations in the FIB that point to a null adjacency are dropped. This is commonly seen in the scenario where a static route pointing to Null0 is configured on an access server being accessed by numerous dialup users. As dialup users are connected to the access server host, routes are installed in the access server's routing table.

An example of a FIB entry pointing to a null adjacency is shown in Example 2-19**.** ICMP unreachable messages are generated for null adjacencies.

Example 2-19 *Null Adjacency*

```
ip route 192.168.1.0 255.255.255.0 Null0
Router#show ip cef 192.168.1.0
192.168.1.0/24, version 37, epoch 0, attached
0 packets, 0 bytes
  via Null0, 0 dependencies
    valid null adjacency
```