CHAPTER 4

Basic IP Connectivity and CEF Troubleshooting

Cisco Express Forwarding (CEF) troubleshooting can be tedious, laborious, and difficult. However, most instances of CEF troubleshooting do not require detailed Cisco IOS architecture and platform (hardware) architecture knowledge. For example, many CEF issues are found in two or three steps of troubleshooting. In addition, many issues that appear to be CEF-related end up being a result of a misconfiguration or inoperable end device.

The first section of this chapter presents the general troubleshooting used on Cisco IOS routers and switches as a first step in troubleshooting IP connectivity problems. CEF occasionally is the scapegoat for IP connectivity problems, and this chapter helps you verify whether CEF is the root cause of a particular IP connectivity problem.

This chapter does not delve into platform specifics of troubleshooting CEF. The chapter simply approaches troubleshooting from a CEF software-switching and command-line interface (CLI) perspective. Most mid- to high-end routers and all Catalyst switches support distributed CEF (dCEF), or hardware switching.

This chapter begins the CEF troubleshooting for all Cisco platforms, including the Cisco 2600, 3700, 7500, 12000, and Catalyst 6500. Chapter 5, "Understanding Packet Switching on the Cisco Catalyst 6500 Supervisor 720," goes into further detail for additional platform and hardware troubleshooting of CEF for the Cisco Catalyst 6500.

The chapter concludes with a table of the basic CEF troubleshooting commands.

Troubleshooting IP Connectivity

As mentioned in the introduction, CEF is a common scapegoat for IP connectivity issues. As such, when approaching an IP connectivity issue, keep an open mind about the root cause of the issue.

This section reviews the methodology for troubleshooting IP connectivity issues, which leads to identifying and troubleshooting CEF issues.

NOTE This chapter is based on Cisco IOS Release 12.3. All the command output presented in this chapter might not be available in previous releases of Cisco IOS. Check the command reference for your specific Cisco IOS version to verify whether a specific command is supported.

The best approach in troubleshooting is to build a troubleshooting plan. Flow charts simplify troubleshooting because they present a stepwise approach to troubleshooting. The following list briefly outlines the first steps in troubleshooting IP connectivity issues and Cisco IOS CEF:

Step 1 Accurately describe the problem.

Step 2 Scope the network topology.

Step 3 Review the Open Systems Interconnection (OSI) model.

(a)Verify the physical layer.

(b)Verify the Layer 2 layer topology.

Step 4 Verify the Address Resolution Protocol (ARP) table.

Step 5 Verify the IP routing table.

Step 6 Verify the CEF Forwarding Information Base (FIB) table.

Step 7 Verify the adjacency table.

Step 8 Conduct hardware-specific troubleshooting.

Accurately Describe the Problem

Accurately articulating your IP connectivity problem is paramount to troubleshooting effectively. An ad hoc approach to troubleshooting is usually ineffective in resolving problems. For example, you do not go to your dentist and tell him you are in pain without describing the symptoms, such as which tooth, how often, how intense, how widespread, what causes the pain, and so on. The same premise exists with IP connectivity troubleshooting. To help yourself, you need to know as much about the issue as possible.

The following questions aid you in accurately articulating your IP connectivity problem:

- Is your IP connectivity problem isolated to a single end device or multiple end devices?
- Is your IP connectivity problem isolated to a single router or Ethernet switch?
- Does your IP connectivity problem exist only on end devices, or does it affect the management CLI of routers and switches as well?
- How widespread is the problem?
- Is the problem widespread or localized to specific area of your network topology?
- Is the problem intermittent or consistent? For example, using the Internet Control Message Protocol (ICMP) ping utility in Cisco IOS and on end devices, are you getting intermittent responses to ICMP echo requests such as every other response, no responses, or inconsistent responses (one out of ten)?

- Does this issue depend on packet size? If you send ICMP echo requests at different sizes, do you consistently get all your responses or does the problem vary with packet size?
- When did the problem first occur? Were there any changes to the network at the same time the problem started occurring?

These questions aid you in articulating your IP connectivity issues. The next section describes an important next step, building the network topology.

Scoping the Network Topology

It is nearly impossible to troubleshoot any type of CEF issue or network connectivity issue without a network diagram that depicts IP addresses, IP routes, devices such as firewalls and switches, and so on. Troubleshooting IP connectivity problems without the aid of a visual topology is nearly impossible unless you can localize the issue to a specific router or switch. In large IP routing scenarios, a network topology is required to troubleshoot connectivity problems. Generally, both logical and physical topologies aid in troubleshooting. Figure 4-1 illustrates a sample physical topology.

Figure 4-1 *Sample Network Topology*

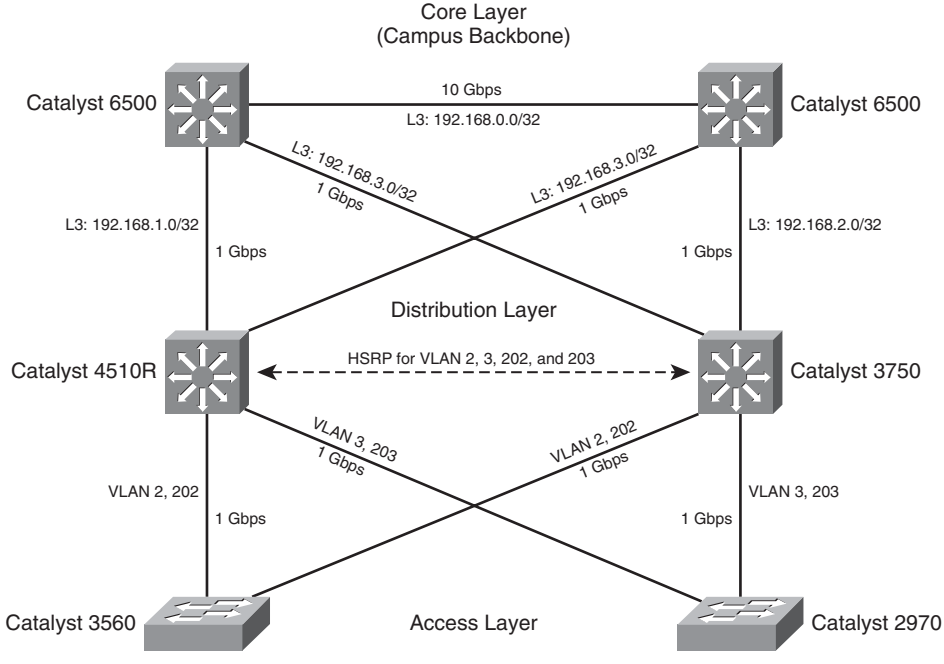

Reviewing the OSI Model for Troubleshooting

The next step in troubleshooting any IP connectivity issue is to review the OSI model and verify that your issue is indeed a Layer 3 (network) issue. Figure 4-2 briefly reviews the OSI model. As a reader of this high-level technology book on CEF, we assume you have an understanding of the OSI model.

Figure 4-2 *OSI Model*

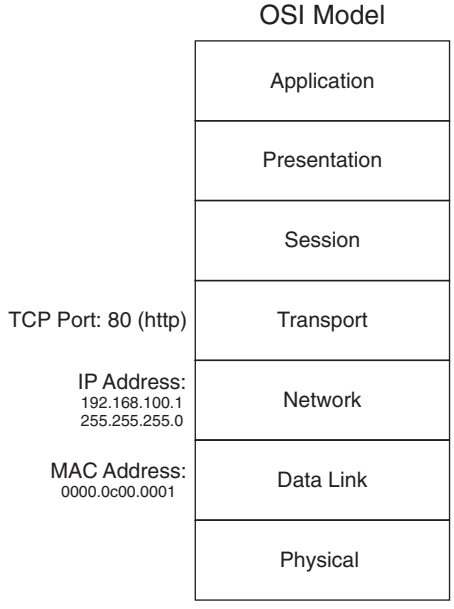

The following sections start from the bottom of the OSI model and review troubleshooting physical connectivity and Layer 2 issues that can affect IP connectivity and give the appearance of a CEF issue.

Troubleshooting Physical Connectivity

An IP connectivity issue might simply be a Layer 1 (physical layer) problem. For example, if you are unable to ping a network device through a router, do not assume that you are having a CEF issue. First, ensure that the host is connected and verify that the physical layer between the host and destination is not sustaining errors. Example 4-1 illustrates sample output from a **show interfaces** command in Cisco IOS.

Example 4-1 *Verifying the Physical Layer*

```
Switch#show interfaces GigabitEthernet 3/6
GigabitEthernet3/6 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet Port, address is 0010.7bfa.808d (bia 0010.7bfa.808d)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is auto, media type is 1000BaseSX
  input flow-control is on, output flow-control is off
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:22, output never, output hang never
  Last clearing of "show interface" counters 5w0d
  Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 2000 bits/sec, 2 packets/sec
     683355349 packets input, 357724908540 bytes, 0 no buffer
     Received 155022 broadcasts (155022 multicast)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 input packets with dribble condition detected
     338647150 packets output, 23295655000 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier
     0 output buffer failures, 0 output buffers swapped out
```

In regard to the **show interfaces** command and verifying physical connectivity, verify that your ingress and egress interfaces are not sustaining errors such as input errors, cyclic redundancy check (CRC) errors, output errors, excessive collisions, overruns, late collisions, or output buffer failures. These types of errors can lead to intermittent or total loss of IP connectivity. Generally, physical layer issues cause intermittent connectivity if the connection has link. Layer 1 errors can be a result of a bad cable, bad port, faulty hardware, and so on. Before proceeding with IP connectivity troubleshooting and ultimately CEF troubleshooting, you must rule out Layer 1 issues.

Troubleshooting Layer 2 Issues

The next step in troubleshooting intermittent connectivity, after physical layer (Layer 1) issues have been ruled out, is to verify that the Layer 2 topology is operating correctly. Verifying the Layer 2 topology includes, but is not limited to, verifying the spanning-tree topology, MAC address table, and Hot Standby Routing Protocol (HSRP) status in switched networks. In a point-to-point router configuration, these features are seldom used. Nevertheless, if your topology involves aggregating and connecting routers to Ethernet switches, you must check the Layer 2 topology before assuming that CEF is the root cause

of an IP connectivity issue. Because CEF is found on all Layer 3 Cisco Catalyst switches, troubleshooting Layer 2 issues is necessary when troubleshooting IP connectivity issues on these platforms.

The following list articulates the most common Layer 2 features that can affect IP connectivity when not operating in a correct state:

- 802.1D, 802.1w, and 802.1s Spanning Tree Protocols
- Hot Standby Routing Protocol (HSRP)
- Virtual Router Redundancy Protocol (VRRP)
- EtherChanneling (port channeling)
- Compression (software- or hardware-based)
- Encryption (software- or hardware-based)
- Firewalls
- Virtual Private LAN Services (VPLS)
- Any Transport over Multiprotocol Label Switching (AToM)

Consult the list for applicable features in your network topology and refer to the Cisco.com website for more details on troubleshooting these features. Troubleshooting these features is outside the scope of this book.

Because many Cisco IOS router experts are not also Cisco IOS switch experts, router experts can find switch troubleshooting daunting. Nevertheless, you should rule out Layer 2 issues as a cause for IP connectivity and CEF issues. The most efficient way to troubleshooting Layer 2 issues is to remove redundant paths, disable as many features as possible, and isolate connectivity issues to a single host.

After the mundane Layer 1 and Layer 2 issues have been ruled out, you can finally transition to troubleshooting IP and CEF, the main focus of this book. As with any troubleshooting, you use a stepwise approach.

Verifying the ARP Table

The first step in troubleshooting IP connectivity issues and CEF from a Layer 3 perspective is to consult the ARP table for identified IP devices that are experiencing connectivity issues. When CEF cannot locate a valid adjacency for a destination prefix, it punts the packets to the CPU for ARP resolution and, in turn, completion of the adjacency.

For example, if the ARP table already lists a particular host, punting it to the process level does not trigger an ARP request. If an entry is incomplete (no response from ARP request) or incorrect in the ARP table, it is also incomplete or incorrect in the CEF adjacency table. This section covers reviewing the ARP table.

Example 4-2 illustrates sample output from the **show arp** command used to display the contents of the ARP table.

Example 4-2 *Displaying the ARP Table in Cisco IOS*

```
Router-2#show arp
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.18.114.250          0   0007.e978.ef03  ARPA   Vlan114
Internet  172.18.114.244        237   0004.7553.cf3c  ARPA   Vlan114
Internet  172.18.114.243          0   Incomplete      ARPA   Vlan114
```

Based on the **show arp** command output, investigate whether the ARP table information is correct. In this example, you see three entries. The 172.18.114.243 entry is incomplete, which means that the device with the address 172.18.114.243 did not respond to the ARP request or is simply powered off. To verify whether the other listings are correct, you might need to access the end device and verify the locally configured MAC address.

Example 4-3 illustrates determining the IP address and MAC address of a Microsoft Windows XP laptop.

Example 4-3 *Determining the IP Address and MAC Address of a Microsoft Windows XP Laptop*

```
C:\WINDOWS\system32>ipconfig.exe /all

Windows IP Configuration

        Host Name . . . . . . . . . . . . : test-winxp
        Primary Dns Suffix  . . . . . . . : amer.bcmsn.com
        Node Type . . . . . . . . . . . . : Hybrid
        IP Routing Enabled. . . . . . . . : No
        WINS Proxy Enabled. . . . . . . . : No
        DNS Suffix Search List. . . . . . : cisco.com
                                            cisco.com

Ethernet adapter Local Area Connection 2:

        Connection-specific DNS Suffix  . : cisco.com
        Description . . . . . . . . . . . : (text deleted) Network Connection
        Physical Address. . . . . . . . . : 00-09-6B-60-15-74
        Dhcp Enabled. . . . . . . . . . . : Yes
        Autoconfiguration Enabled . . . . : Yes
        IP Address. . . . . . . . . . . . : 172.18.114.244
        Subnet Mask . . . . . . . . . . . : 255.255.0.0
        Default Gateway . . . . . . . . . : 172.18.114.1
        DHCP Server . . . . . . . . . . . : 10.18.0.1
        DNS Servers . . . . . . . . . . . : 10.2.2.1
        Lease Obtained. . . . . . . . . . : Saturday, February 12, 2005 12:08:5 PM
        Lease Expires . . . . . . . . . . : Sunday, February 13, 2005 12:38:59 PM

Ethernet adapter Wireless Network Connection 3:
```

continues

Example 4-3 *Determining the IP Address and MAC Address of a Microsoft Windows XP Laptop (Continued)*

```
        Media State . . . . . . . . . . . : Media disconnected
        Description . . . . . . . . . . . : Cisco Systems PCI Wireless LAN Adapter
        Physical Address. . . . . . . . . : 00-02-8A-3E-6D-CB
```

In Example 4-3, the MAC (physical) address of the host is 00-09-6B-60-15-74. This differs from the ARP table of the local router. A protocol analyzer such as Ethereal is necessary to decode frames on Ethernet to investigate why the router has an incorrect entry. Example 4-4 shows an example of troubleshooting ARP entries with the **debug arp** command.

NOTE Because this is simply a mock setup for illustrative purposes, using a Cisco IOS debug is not an issue. Generally, in large networks, debugs such as **debug arp** are intrusive and should not be used in production environments.

Example 4-4 *Troubleshooting Incorrect ARP Entries Using the* **debug arp** *Command*

```
Router-2#debug arp
ARP packet debugging is on
Router-2#clear arp int vlan 114
*Feb 12 10:43:23.710 UTC: IP ARP: sent req src 172.18.114.4 0008.a378.bdff,
                 dst 172.18.114.250 0007.e978.ef03 Vlan114
*Feb 12 10:43:23.710 UTC: IP ARP: sent req src 172.18.114.4 0008.a378.bdff,
                 dst 172.18.114.244 0004.7553.cf3c Vlan114
*Feb 12 10:43:23.718 UTC: IP ARP: rcvd rep src 172.18.114.250 0007.e978.ef03, dst
  172.18.114.4 Vlan114
*Feb 12 10:43:23.718 UTC: IP ARP: creating entry for IP address: 172.18.114.250,
  hw: 0007.e978.ef03
*Feb 12 10:43:23.718 UTC: IP ARP: rcvd rep src 172.18.114.244 0009.6B60.1574, dst
  172.18.114.4 Vlan114
*Feb 12 10:43:23.718 UTC: IP ARP: creating entry for IP address: 172.18.114.244,
  hw: 0009.6B60.1574
*Feb 12 10:43:24.124 UTC: IP ARP: rcvd rep src 172.18.114.244 0004.7553.cf3c, dst
  172.18.114.4 Vlan114
*Feb 12 10:43:24.124 UTC: IP ARP: creating entry for IP address: 172.18.114.244,
  hw: 0004.7553.cf3c
```

From the debug output, it is apparent that two devices are replying to the ARP request sent by the router: the test host and a rogue or misconfigured device. The router always populates the ARP table with the most recent ARP response. Because the rogue entry arrived second, the ARP table was incorrect. Another example of invalid ARP entries is with Frame Relay point-to-point or other types of point-to-point interfaces. Because the ARP entry is incorrect, the CEF adjacency table entry will also be incorrect.

NOTE When the ping process on a Cisco router or switch running Cisco IOS attempts to send an ICMP echo to a host for which an ARP entry does not exist, the router or switch initiates an ARP request. Because the first ICMP echo is dropped after a defined period awaiting the ARP response, the first ICMP echo fails to be sent. As a result, initiating an ICMP ping with default parameters to a device for which a current ARP entry does not exist generally results in a success rate of 4/5 (80 percent) because five ICMP echoes are sent by default in Cisco IOS.

Moreover, a common IP connectivity problem associated with CEF is where a more specific route, such as a host route (/32), is learned on a different interface than the connected route (for example, a host, 192.168.1.10, connected on interface Ethernet 0/0 of Cisco IOS router). According to the routing table, 192.168.1.0/24 is learned as a directly connected route. If the router learns a more specific route, such as 192.168.1.10/32, from another interface, the router forwards the packets to the more specific route. This is a common issue; you should rule it out early in troubleshooting.

Verifying the Routing Table

The example in the previous section was based on a directly connected device. To illustrate verifying the IP routing table and subsequent CEF troubleshooting, Example 4-5 is based on a router with several routes, as illustrated in Figure 4-3. Devices connected to Router 2 in the 10.1.1.0/24 subnet are unable to reach a Secure File Transfer Protocol (SFTP) server with the address 10.18.118.184, as shown in the first steps of Example 4-5; specifically, the host 10.1.1.100 is unable to ping 172.18.118.184. For the purpose of this example, assume that all the IP routing configuration and host configurations, such as IP address and default gateways, are configured correctly. In addition, assume that our issue is strictly limited to Router 2 and not Router 1, the Layer 3 switch.

Figure 4-3 *Troubleshooting IP Connectivity and Verifying the IP Routing Table*

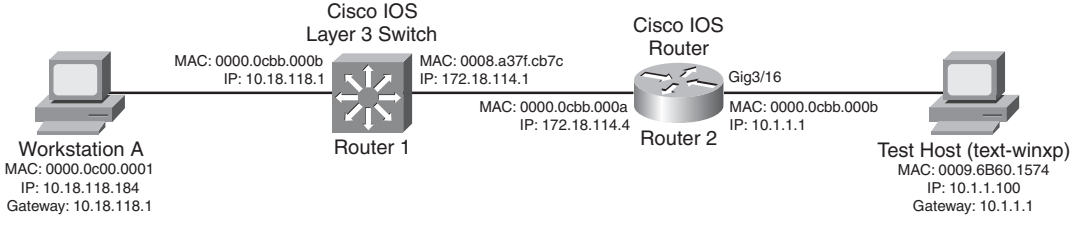

Example 4-5 *IP Connectivity Issue from the Host, 10.1.1.100, to the Target, 172.18.118.184, Based on Figure 4-3*

```
Host 10.1.1.100:

C:\WINDOWS\system32>ping 10.18.118.184

Pinging 172.18.118.184 with 32 bytes of data:

Request timed out.
Request timed out.
Request timed out.
Request timed out.

Ping statistics for 172.18.118.184:
    Packets: Sent = 4, Received = 0, Lost = 4 (100% loss),

C:\WINDOWS\system32>ipconfig

Windows IP Configuration


Ethernet adapter Local Area Connection 2:

        Connection-specific DNS Suffix  . : cisco.com
        IP Address. . . . . . . . . . . . : 10.1.1.100
        Subnet Mask . . . . . . . . . . . : 255.255.255.0
        Default Gateway . . . . . . . . . : 10.1.1.1

C:\WINDOWS\system32>tracert 10.18.118.183

Tracing route to 10.18.118.183 over a maximum of 30 hops

  1   <10 ms   <10 ms   <10 ms  Router-2.cisco.com [10.1.1.1]
  2     *        *        *     Request timed out.
  3     *        *        *     Request timed out.
  4     *        *        *     Request timed out.
  5     *        *        *     Request timed out.
  6     *        *        *     Request timed out.
!Output omitted
```

The first step in troubleshooting is to access the first-hop router, Router 2, and verify IP connectivity to the SFTP server, 10.18.118.184. Example 4-6 also confirms the route to the host, 10.18.118.184, and verifies the ARP entry for the next hop for 10.18.118.184, which is 10.18.114.1. The host entry, 10.18.118.184, is known through a static route with a metric of 0 according to the **show ip route** command. This is the gateway of last resort.

Example 4-6 *Verifying IP Routing Table and Next-Hop Information*

```
Router-2#ping 172.18.118.184

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.18.118.184, timeout is 2 seconds:
!!!!!
```

Example 4-6 *Verifying IP Routing Table and Next-Hop Information (Continued)*

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
Router-2#

Router-2#show ip route 172.18.118.184
Routing entry for 172.18.118.0/24
  Known via "static", distance 1, metric 0
  Routing Descriptor Blocks:
  * 172.18.114.1
      Route metric is 0, traffic share count is 1

Router-2#show ip arp 172.18.114.1
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  172.18.114.1            0   0008.a37f.cb7c  ARPA   Vlan114

Router-2#show ip route 172.18.114.1
Routing entry for 172.18.114.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via ospf 1
  Routing Descriptor Blocks:
  * directly connected, via Ethernet0/0
      Route metric is 0, traffic share count is 1
```

Although the ping from the router was successful, the next step is to verify the MAC (Hardware Addr) of the next-hop router. This step verifies that the correct ARP entry exists for the next-hop router. Example 4-7 illustrates obtaining Router 1's MAC address for the interface configured with the IP address 172.18.114.1.

Example 4-7 *Verifying the MAC Address of an Interface in Cisco IOS*

```
Router-1#show interfaces vlan 114 ¦ include address
!Output omitted
  Hardware is Cat6k RP Virtual Ethernet, address is 0008.a37f.cb7c (bia
0008.a37f.cb7c)
  Internet address is 172.18.114.1/24
!Output omitted
```

The next-hop router's (Router 1's) MAC address is indeed 0008.a37f.cb7c. Therefore, at this point in troubleshooting, both the routing entry and ARP entry for the next-hop router are correct.

The next step is to verify IP connectivity to the next-hop router, Router 1, from both the host, 10.1.1.100, and Router 2, as shown in Example 4-8.

Example 4-8 *Verifying IP Connectivity to the Next Hop from Router 2 and the Host, 10.1.1.100*

```
Router-2#ping 172.18.114.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.18.114.1, timeout is 2 seconds:
!!!!!
```

continues

Example 4-8 *Verifying IP Connectivity to the Next Hop from Router 2 and the Host, 10.1.1.100 (Continued)*

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

Router-2#ping
Protocol [ip]:
Target IP address: 172.18.114.1
Repeat count [5]: 10
Datagram size [100]: 1500
Timeout in seconds [2]:
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 10, 1500-byte ICMP Echos to 172.18.114.1, timeout is 2 seconds:
!!!!!!!!!!
Success rate is 100 percent (10/10), round-trip min/avg/max = 1/2/4 ms

Host 10.1.1.100:

C:\WINDOWS\system32>ping 172.18.114.1

Pinging 172.18.118.184 with 32 bytes of data:

Reply from 172.18.114.1: bytes=32 time=41ms TTL=253
Reply from 172.18.114.1: bytes=32 time=41ms TTL=253
Reply from 172.18.114.1: bytes=32 time=41ms TTL=253
Reply from 172.18.114.1: bytes=32 time=40ms TTL=253

Ping statistics for 172.18.114.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 40ms, Maximum = 41ms, Average = 40ms
```

Example 4-8 also issued an ICMP ping of a larger packet size for continued verification of accessibility of the next-hop router. You must verify various packet sizes because default-sized packets of ICMP echoes can be routed and switched correctly in the network, whereas larger packets can be dropped for a multitude of reasons. These reasons include encoding errors in WAN circuits, fragmentation configuration issues, hardware issues, Virtual Private Network (VPN) misconfiguration, software features such as encryption and compression, and so on. Some network routers' security policies can also prevent the routers from responding to very high rates of ICMP packets. Therefore, you should understand your network policies when troubleshooting or installing systems. In Example 4-8, the router is indeed reachable through ICMP.

Example 4-8 proved that the next hop is reachable from both Router 2 and the host, 10.1.1.100. Therefore, at this point, the ARP and IP routing tables appear to be correct on Router 2. The next step is troubleshooting CEF (Step 6), which is discussed in the next section using the same example as that in Figure 4-3.

Sending ICMP echoes, requests, or responses from the CLI of the router is software switched, while ICMP echoes from end devices through many Cisco IOS routers and all current-generation Catalyst switches are hardware switched. The software-switching path might be correct, but the hardware-switching path might not be correct. In this example, the software-switching path on Router 2 might be correct, but the hardware-switching path might not be correct. Troubleshooting the hardware-switching path is outside the scope of this chapter. Chapter 5 discusses troubleshooting hardware-switching paths on a Cisco Catalyst 6500 platform. However, you must troubleshoot CEF from the software-switching perspective first because the hardware-switching tables are built from the software-switching tables.

The section "Troubleshooting the CEF FIB Table," later in this chapter, continues the investigation of the software-switching path by troubleshooting the Cisco IOS CEF table on Router 2 for Figure 4-3. As noted previously, because the hardware-switching path gets built from the software-switching CEF and adjacency table, you should investigate the software-switching path first.

Using IOS Ping with the Record Option to Rule Out CEF

CEF does not support all IP packet types and must process-switch specific types of packets. One such packet includes the ICMP echo with the record option. As such, you can rule out CEF as a cause of an IP connectivity problem with some certainty using the ICMP echo with record option.

For example, in Figure 4-3, sending an ICMP echo with the record option forces all routers along the path to use the process-switching method of forwarding a frame. If an ICMP echo with the record option is successful and a standard ICMP echo is not, you can assume with some certainty that CEF is indeed a cause of your IP connectivity issue somewhere along the path.

In Cisco IOS, use the **ping** command with extended commands option to send ICMP echoes with the record option. Example 4-9 illustrates the use of the ICMP echo with record option.

Example 4-9 *Sending ICMP Packets Using the Record Option*

```
Router-2#ping ip
Target IP address: 10.18.118.184
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: loop 0
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]: r
```

continues

Example 4-9 *Sending ICMP Packets Using the Record Option (Continued)*

```
Number of hops [ 9 ]:
Loose, Strict, Record, Timestamp, Verbose[RV]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.18.118.184, timeout is 2 seconds:
Packet has IP options:  Total option bytes= 39, padded length=40
 Record route: <*>
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)

Reply to request 0 (1 ms).  Received packet has options
 Total option bytes= 40, padded length=40
 Record route:
   (172.18.114.4)
   (10.18.118.1)
   (10.18.118.184)
   (172.18.114.4) <*>
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
   (0.0.0.0)
 End of list

!Output omitted for brevity
```

NOTE When troubleshooting CEF, always note that certain packets, mainly those requiring special handling, are not CEF switched.

Troubleshooting the CEF FIB Table

The following sections outline basic CEF troubleshooting using a stepwise approach based on Figure 4-3 and Example 4-5. These sections move to focusing on CEF.

The steps for verifying the CEF table are as follows:

Step 1 Verify the CEF configuration.

Step 2 Confirm the IP CEF switching path, including using CEF accounting counters to confirm the switching path.

Step 3 Verify CEF switching details.

Verifying the CEF Configuration

When verifying the CEF table (FIB), first verify whether CEF is indeed enabled globally and on an interface basis using the following commands:

```
show ip cef
show cef interface <module_number>/<port_number>
```

Example 4-10 illustrates how to verify that CEF is indeed enabled on both a global and an interface basis.

Example 4-10 *Verifying That IP CEF Is Enabled Globally and Per-Interface*

```
Router-2#show ip cef
Prefix              Next Hop            Interface
0.0.0.0/0           172.18.114.1        Ethernet0/0
0.0.0.0/32          receive
10.0.0.0/24         172.18.114.1        Ethernet0/0
10.1.1.0/24         attached            Ethernet0/1
10.1.1.0/32         receive
10.1.1.1/32         receive
10.1.1.100/32       10.1.1.100          Ethernet0/1
10.1.1.255/32       receive
10.18.118.0/24      172.18.114.1        Ethernet0/0
10.224.0.0/24       172.18.114.1        Ethernet0/0
10.225.0.0/24       172.18.114.1        Ethernet0/0
10.226.0.0/24       172.18.114.1        Ethernet0/0
165.27.1.0/24       172.18.114.1        Ethernet0/0
172.18.114.0/24     attached            Ethernet0/0
172.18.114.0/32     receive
172.18.114.1/32     172.18.114.1        Ethernet0/0
172.18.114.4/32     receive
172.18.114.5/32     172.18.114.5        Ethernet0/0
172.18.114.7/32     172.18.114.7        Ethernet0/0
172.18.114.177/32   172.18.114.177      Ethernet0/0
172.18.114.191/32   172.18.114.191      Ethernet0/0
172.18.114.214/32   172.18.114.214      Ethernet0/0
Prefix              Next Hop            Interface
172.18.114.255/32   receive
172.18.116.64/29    172.18.114.1        Ethernet0/0
192.168.100.0/24    172.18.114.1        Ethernet0/0
224.0.0.0/4         drop
224.0.0.0/24        receive
255.255.255.255/32  receive
```

If CEF is not enabled globally, the **show ip cef** command returns the message "%CEF not running," as shown in Example 4-11.

Example 4-11 *Example of Router Not Running CEF*

```
Router-2#show ip cef

%CEF not running
```

Although the previous step shows that CEF is enabled on the interfaces in respect to Figure 4-3, the **show ip interface** and the **show cef interface** *interface* commands confirm that CEF is enabled on a per-interface basis. Example 4-12 illustrates the **show ip interface** command.

Example 4-12 *Verifying That CEF Is Enabled on a Per-Interface Basis Using the* **show ip interface** *Command*

```
Router-2#show ip interface ethernet 0/0
Ethernet0/0 is up, line protocol is up
  Internet address is 172.18.114.4/24
  Broadcast address is 255.255.255.255
  Address determined by setup command
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Multicast reserved groups joined: 224.0.0.5 224.0.0.6
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is enabled
  IP fast switching on the same interface is disabled
  IP Flow switching is disabled
  IP CEF switching is enabled
  IP CEF Fast switching turbo vector
  IP multicast fast switching is enabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are Fast, CEF
  Router Discovery is disabled
  IP output packet accounting is disabled
  IP access violation accounting is disabled
  TCP/IP header compression is disabled
  RTP/IP header compression is disabled
  Probe proxy name replies are disabled
  Policy routing is disabled
  Network address translation is disabled
  Web Cache Redirect is disabled
  BGP Policy Mapping is disabled

Router-2#show cef interface ethernet 0/0
Ethernet0/0 is up (if_number 2)
  Corresponding hwidb fast_if_number 2
```

Example 4-12 *Verifying That CEF Is Enabled on a Per-Interface Basis Using the* **show ip interface** *Command (Continued)*

```
  Corresponding hwidb firstsw->if_number 2
  Internet address is 172.18.114.4/24
  ICMP redirects are always sent
  Per packet load-sharing is disabled
  IP unicast RPF check is disabled
  Inbound access list is not set
  Outbound access list is not set
  IP policy routing is disabled
  BGP based policy accounting is disabled
  Hardware idb is Ethernet0/0
  Fast switching type 1, interface type 61
  IP CEF switching enabled
  IP CEF Feature Fast switching turbo vector
  Input fast flags 0x0, Output fast flags 0x0
  ifindex 1(1)
  Slot 0 Slot unit 0 Unit 0 VC -1
  Transmit limit accumulator 0x0 (0x0)
  IP MTU 1500
```

Confirming the IP CEF Switching Path

The next step in troubleshooting the CEF table is to verify the switching path that the router in question is using. The **show interfaces stat** command displays the switching path stats on Cisco IOS routers, as shown in Example 4-13. The processor row includes process-switched (software-switched) packets, while the router cache row includes both CEF-switched and fast-switched packets. Note that high-end routers and switches that support distributed CEF display an additional row referred to as Hardware, Parallel Express Forwarding (PXF), or Distributed.

Example 4-13 *Displaying Interface Switching Statistics*

```
Router-2#show interfaces stat
Ethernet0/0
          Switching path    Pkts In   Chars In   Pkts Out  Chars Out
               Processor     398302   28590641      34250    3035319
            Route cache        173      81340        192      12152
                  Total     398475   28671981      34442    3047471


Ethernet0/1
          Switching path    Pkts In   Chars In   Pkts Out  Chars Out
               Processor      14614    3087533      18038    1700974
            Route cache        191      12955        172      80920
                  Total      14805    3100488      18210    1781894
```

From Example 4-13, the router is using route cache for switching. The command does not tell you how much is CEF switched or how much is fast switched because route cache

includes both CEF-switched and fast-switched packets. Another way to determine whether the router is fast-switching or CEF-switching packets is to view the contents of the fast-switching table using the **show ip cache** command. If prefix entries exist in the fast-switching table, the traffic to the destination entries are being fast switched. Example 4-14 illustrates using the **show ip cache** command to verify CEF switching.

Example 4-14 *Viewing the Fast-Switching (IP Cache) Table*

```
Router-2#show interface stat
Ethernet0/0
          Switching path    Pkts In   Chars In   Pkts Out  Chars Out
               Processor     400445   28741844      34413    3049576
            Route cache        416     128348        436      29516
                  Total     400861   28870192      34849    3079092
Interface Serial0/0 is disabled

Ethernet0/1
          Switching path    Pkts In   Chars In   Pkts Out  Chars Out
               Processor      14662    3099087      18119    1708563
            Route cache        436      30581        416     127949
                  Total      15098    3129668      18535    1836512

Router-2#show ip cache
IP routing cache 0 entries, 0 bytes
   2 adds, 2 invalidates, 0 refcounts
Minimum invalidation interval 2 seconds, maximum interval 5 seconds,
   quiet interval 3 seconds, threshold 0 requests
Invalidation rate 0 in last second, 0 in last 3 seconds
Last full cache invalidation occurred 1d18h ago

Prefix/Length          Age       Interface       Next Hop

Router-2#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router-2(config)#no ip cef
Router-2(config)#end

Router-2#show ip cef
%CEF not running
Prefix              Next Hop             Interface

Router-2#show ip cache
IP routing cache 2 entries, 344 bytes
   4 adds, 2 invalidates, 0 refcounts
Minimum invalidation interval 2 seconds, maximum interval 5 seconds,
   quiet interval 3 seconds, threshold 0 requests
Invalidation rate 0 in last second, 0 in last 3 seconds
Last full cache invalidation occurred 1d18h ago

Prefix/Length          Age       Interface       Next Hop
10.1.1.100/32           00:00:10  Ethernet0/1     10.1.1.100
10.18.118.0/24          00:00:10  Ethernet0/0     172.18.114.1
```

Example 4-14 *Viewing the Fast-Switching (IP Cache) Table (Continued)*

```
Router-2#show interfaces stat
Ethernet0/0
          Switching path    Pkts In   Chars In   Pkts Out  Chars Out
               Processor     400665   28758026      34436    3051672
            Route cache        506     135008        525      36102
                  Total     401171   28893034      34961    3087774


Ethernet0/1
          Switching path    Pkts In   Chars In   Pkts Out  Chars Out
               Processor      14674    3101447      18131    1709787
            Route cache        524      37093        505     134535
                  Total      15198    3138540      18636    1844322

Router-2#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router-2(config)#ip cef
Router-2(config)#end

Router-2#show ip cache
IP routing cache 0 entries, 0 bytes
   5 adds, 5 invalidates, 0 refcounts
Minimum invalidation interval 2 seconds, maximum interval 5 seconds,
   quiet interval 3 seconds, threshold 0 requests
Invalidation rate 0 in last second, 0 in last 3 seconds
Last full cache invalidation occurred 1d18h ago

Prefix/Length          Age       Interface       Next Hop

Router-2#
```

Example 4-14 illustrates that the route cache counters from the **show interface stats** command were CEF switched by the router in our example. Example 4-14 views the output of the **show interface stats** command, checks the fast-switching table using the **show ip cache** command, and then disables CEF. After disabling CEF, the fast-switching table was rechecked and entries were observed.

In reference to Figure 4-3, where the host 10.1.1.100 could not ping 10.18.118.184, the troubleshooting step of ruling in or ruling out CEF is to disable CEF switching, globally or per interface, and reattempt to send the ICMP echoes. Of course, in a production environment, disabling CEF might not be an option; however, during a change-control window or isolated situation, this troubleshooting is useful. Always consult with a Cisco Technical Assistance Center (TAC) engineer before disabling CEF.

When disabling CEF on a per-interface basis for troubleshooting, pay close attention to the ingress and egress interfaces. Generally, when disabling CEF on a per-interface basis for troubleshooting, disable CEF on both the ingress and egress interface. In a mixed-mode

environment, where the ingress and egress interface use different switching configurations, refer to Table 4-1 for the resulting switching method.

Table 4-1 *Resulting Switching Method Based on Ingress and Egress Switching Configuration*

Ingress Interface	Egress Interface	Resulting Switching Method
CEF	Process	CEF
Process	CEF	Fast
Process	Fast Switching	Fast Switching
CEF	Fast Switching	CEF

Based on the information in Table 4-1, CEF switching occurs on the ingress. Therefore, use the **no ip route-cache cef** command on the ingress interface to disable CEF. In contrast, because Cisco IOS builds a fast-switching cache entry after switching a packet, packets ingress on a process-switched interface and egress through a fast-switched interface. Therefore, use the **no ip route-cache** command on the egress interface to disable fast switching. Alternatively, disabling CEF on a global basis is permissible on several platforms, mostly low-end platforms that use the **no ip cef** command. Check the release notes and configuration guide for your specific platform when attempting to disable CEF globally or per interface.

NOTE To reenable CEF on a per-interface basis, both fast switching and CEF have to be enabled. As a result, not only is the **ip route-cache cef** interface configuration command required, but the **ip route-cache** interface configuration command is also required for CEF to function on a per-interface basis.

Furthermore, newer mid- to high-end platforms do not support disabling CEF on a global or per-interface basis. When attempting to disable CEF on a platform such as the Catalyst 6500, Cisco 12000, CRS, and Cisco 7600 that do not support disabling CEF, the CLI returns an error. Example 4-15 illustrates an example of attempting to disable CEF on a platform that does not support disabling CEF.

Example 4-15 *Attempting to Disable CEF on a Platform Such as the Catalyst 4500 That Does Not Support Disabling CEF*

```
Catalyst4500#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Catalyst4500(config)#no ip cef
%Cannot disable CEF on this platform
```

Using CEF Accounting Counters to Confirm the Switching Path

Cisco IOS supports CEF accounting, with specific limitations pertaining to hardware-switched traffic. However, with software switching, CEF supports an accounting option for packet and byte counters.

When troubleshooting CEF, you can view the packet and byte counters on a per-prefix basis. To enable CEF accounting, use the **ip cef accounting per-prefix** global configuration command**.** Simply use the **show ip cef** *prefix* command to display the counters for a specific prefix. Example 4-16 illustrates the use of CEF accounting.

Example 4-16 *CEF Accounting*

```
Router-2#configure terminal
Router-2(config)#ip cef accounting per-prefix
Router-2(config)#end
Router-2#show ip cef 10.18.118.84
10.18.118.0/24, version 53, epoch 0, cached adjacency 172.18.114.1
240 packets, 14400 bytes
  via 172.18.114.1, Ethernet0/0, 0 dependencies
    next hop 172.18.114.1, Ethernet0/0
    valid cached adjacency
```

NOTE When you enable network accounting for dCEF from global configuration mode, accounting information grouped by IP prefix is not sent to the route processor (RP) for viewing through the **show ip cef** command. However, the accounting information is collected by dCEF processes on the line card. In this situation, use the **show cef linecard** command to view dCEF statistics.

Verifying the CEF Switching Details

In most production environments, you usually cannot disable CEF, even on a per-interface basis, during normal production. This is because in many production networks, traffic rates exceed the software-switching capabilities of the router, and disabling CEF forces software switching. Therefore, disabling CEF to troubleshoot CEF is not always an option. The next step in these situations is to verify the CEF table information.

Example 4-17 illustrates the use of the **show ip cef** command to gather details about the CEF table on a per-entry basis. The values chosen for this example match those for the troubleshooting example.

Example 4-17 *Gathering CEF Table Details*

```
Router-2#show ip cef 10.18.118.184 detail
10.18.118.0/24, version 23, cached adjacency 172.18.114.1
0 packets, 0 bytes
  via 172.18.114.1, Ethernet0/0, 0 dependencies
```

continues

Example 4-17 *Gathering CEF Table Details (Continued)*

```
    next hop 172.18.114.1, Ethernet0/0
    valid cached adjacency

Router-2#show ip cef 172.18.114.1 detail
172.18.114.1/32, version 17, cached adjacency 172.18.114.1
0 packets, 0 bytes
  via 172.18.114.1, Ethernet0/0, 1 dependency
    next hop 172.18.114.1, Ethernet0/0
    valid cached adjacency
```

From the output shown in Example 4-17, the host 10.18.118.184 next hop is 172.18.114.1 and is a valid cached adjacency through Ethernet0/0. The first step is to verify this output against Figure 4-3 and the **show ip route** command from Example 4-5. The details are correct. The next step is to verify the next-hop CEF entry for 10.18.118.184, which is 172.18.114.1. Again, this information appears to be correct based on Figure 4-3. If this information was not correct, the next step is to either disable CEF and test connectivity or open a Cisco TAC case.

Another possible reason for packet loss with CEF is CEF drop adjacency. CEF drop adjacencies define hardware-switched drops for prefixes. CEF drop adjacencies allow dropping frames in hardware rather than punting every frame for software switching. This is an effective method in preventing denial of service (DoS) attacks and high CPU usage due the CPU processing excessive drops. If a CEF drop adjacency exists, it is generally because of one of the following reasons:

- Unsupported features.
- Packets destined to prefixes associated with punt adjacencies that exceed a predefined rate. Punting rate limiters exist to prevent DoS attacks and high CPU usage caused by excessive punts.
- Unresolved or missing FIB entry.
- Unsupported frame type.

Example 4-18 illustrates a sample output from the **show cef drop** command. In respect to Figure 4-3 and Example 4-5, attempting to relate the CEF drop counters to the ICMP packets can yield additional details about why the packets are being dropped. However, in a production environment, it is difficult to correlate the ICMP packet loss to **show cef drop** counters because production environments generally have multiple flows passing traffic simultaneously. Nonetheless, if the **show cef drop** counters remain 0 during the ping tests, you can rule out the notion that the ping failed because of CEF drop adjacencies. In later code, the **show ip cef switching statistics** command gives detailed information about why a drop occurs and replaces the **show cef not** command.

Example 4-18 **show cef drop** *Command Example*

```
Router-2#show cef drop
CEF Drop Statistics
Slot  Encap_fail  Unresolved Unsupported    No_route      No_adj  ChkSum_Err
RP             3           0           0           0           0           0
```

Table 4-2 defines the fields associated with the **show cef drop** command.

Table 4-2 **show cef drop** *Field Descriptions*

Field	Description
Slot	Refers to the slot for the respective ingress packet counts. For Cisco IOS routers that do not support dCEF, this value is always RP.
Encap_fail	Indicates the number of packets dropped after exceeding the limit for packets punted to the processor because of missing adjacency information such as an unresolved ARP request. Note that CEF throttles packets punted to the process level at a rate of one packet per second to aid in susceptibility of DoS attacks.
Unresolved	Indicates the number of packets dropped because of an unresolved prefix in the FIB table.
Unsupported	Indicates the number of packets fast-dropped by CEF (drop adjacency) because of an unsupported feature.
No_route	Indicates the number of packets dropped because of a missing prefix in the FIB table.
No_adj	Indicates the number of packets dropped because of incomplete adjacency.
Chksum_Err	Indicates the number of IP version 4 (IPv4) packets received with a checksum error.

In controlled environments, preferably in nonproduction environments, use debugs to troubleshoot CEF. If you can narrow your issue to a specific host or subnet and the router or switch under investigation is logging drops or receives, you can use debugs that are limited to specific destinations to troubleshoot the issue. Limiting debugs is done by limiting debug output to specific IP sources or destinations configured in an access list. Example 4-19 illustrates an example of configuring an access control list (ACL) to limit the output of a CEF debug to a specific destination.

Example 4-19 *Controlling Debug Output Defined by an ACL*

```
Router-2#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router-2(config)#access-list 10 permit 10.1.1.1
Router-2(config)#end

Router-2#debug ip cef drop 10
IP CEF drops debugging is on for access list 10
```

continues

Example 4-19 *Controlling Debug Output Defined by an ACL (Continued)*

```
Router-2#debug ip cef receive 10
IP CEF received packets debugging is on for access list 10

*Mar  7 23:27:29.681 UTC: CEF-receive: Receive packet for 10.1.1.1
*Mar  7 23:27:29.681 UTC: CEF-Receive: Packet for 10.1.1.1 -- receive
*Mar  7 23:27:29.681 UTC: IP-CEF: Receive packet for 10.1.1.1 (process switch)
```

Based on the example in Figure 4-3 and Example 4-5, the debug output illustrates the router processing "receive" packets destined for its own interface. Receive packets include packets punted by CEF for software switching. Drop packets appear in the debug output in the same manner. As a result, this debug, used in controlled environments, can yield additional information helpful in troubleshooting.

Verifying the Adjacency Table

After verifying the CEF FIB table, the next step in troubleshooting CEF is to verify the adjacency table. The adjacency table contains the rewrite information that CEF uses to switch packets. Verifying that the rewrite information is correct is an important step in troubleshooting CEF operation.

The four commands that provide different hierarchical levels of information for the adjacency table are **show adjacency**, **show adjacency summary**, **show adjacency detail**, and **show adjacency internal**. Example 4-20 shows examples of these commands, respectively.

Example 4-20 *Viewing Adjacency Table Details*

```
Router-2#show adjacency
Protocol Interface                 Address
IP       Ethernet0/1               10.1.1.100(5)
!Output omitted for brevity
IP       Ethernet0/0               172.18.114.1(23)

Router-2#show adjacency summary
Adjacency Table has 6 adjacencies
  Table epoch: 0 (6 entries at this epoch)

  Interface                Adjacency Count
  Ethernet0/0              5
  Ethernet0/1              1

Router-2#show adjacency detail
Protocol Interface                 Address
!Output omitted for brevity
IP       Ethernet0/1               10.1.1.100(5)
                                   561 packets, 41514 bytes
                                   0007E905156C00503EFA37810800
                                   ARP        03:53:37
```

Example 4-20 *Viewing Adjacency Table Details (Continued)*

```
                                    Epoch: 0
!Output omitted for brevity
IP       Ethernet0/0                172.18.114.1(23)
                                    581 packets, 43012 bytes
Protocol Interface                  Address
                                    0008A37FCB7C00503EFA37800800
                                    ARP        04:02:59
                                    Epoch: 0

Router-2#show adjacency internal
Protocol Interface                  Address
!Output omitted for brevity
IP       Ethernet0/0                172.18.114.1(23)
                                    280 packets, 20738 bytes
                                    0008A37FCB7C00503EFA37800800
                                    ARP        04:02:58
                                    Epoch: 0
                                    Fast adjacency disabled
                                    IP redirect enabled
                                    IP mtu 1500 (0x0)
                                    Fixup disabled
                                    Adjacency pointer 0x816BFC20, refCount 23
                                    Connection Id 0x000000
                                    Bucket 205
IP       Ethernet0/1                10.1.1.100(5)
                                    269 packets, 19906 bytes
                                    0007E905156C00503EFA37810800
                                    ARP        03:58:29
                                    Epoch: 0
                                    Fast adjacency disabled
                                    IP redirect enabled
                                    IP mtu 1500 (0x0)
                                    Fixup disabled
                                    Adjacency pointer 0x816BFAD0, refCount 5
                                    Connection Id 0x000000
                                    Bucket 110
!Output omitted for brevity
```

Table 4-3 describes the most significant fields from the **show adjacency** commands in Example 4-20 for the purpose of troubleshooting. Verifying the information against the **show ip route** and **show arp** commands is necessary in verifying CEF consistency. If the values are not correct, disable CEF as a workaround and open a Cisco TAC case.

Table 4-3 **show adjacency** *Command Field Descriptions*

Field	Description
172.18.114.1(23)	The value in parentheses, 23, refers to the number of times a FIB entry points to an adjacency entry (refCount). Numerous system entries are not shown in the output of the **show ip cef** command. In the hardware used for this example, a minimum of five references per IP address existed. As a result, four additional FIB routing entries point to the IP address.
0008A37FCB7C0008A 378BDFF0800	The first 12 characters, 0008A37FCB7C, are the MAC address associated with the destination next-hop interface (destination MAC address rewrite). The next 12 characters represent the MAC address of the source interface of the packet (source MAC address rewrite). The last four characters represent the well-known Ethertype value 0x0800 for IP for Advanced Research Projects Agency (ARPA) encapsulation, the default among Ethernet vendors.
ARP 04:02:58	The ARP value indicates that the entry was learned through the ARP process. The timestamp indicates the time remaining before the entry times out. The default ARP timeout is 4 hours.
Fast adjacency disabled	An FIB entry caches an adjacency for a next-hop interface when not doing load sharing over multiple active paths. A fast adjacency increases the switching speed of packets.

In Figure 4-3 and Example 4-5, the destination MAC address rewrite information from the **show adjacency detail** command in Example 4-20, 0008A37FCB7C, must match the MAC address from the **show arp** command. Otherwise, if the MAC address did not match, an inconsistency issue exists between the ARP table and adjacency table that needs to be investigated with the Cisco TAC.

Hardware-Specific Troubleshooting

If you are troubleshooting CEF on a Cisco IOS router that is only performing software switching and are unable to find any issues based on the previous sections, the next step is to open a Cisco TAC case and include relevant information. However, if you are troubleshooting a platform that supports dCEF or hardware switching, more analysis is needed.

Chapter 5 provides information on troubleshooting hardware switching on a Cisco Catalyst 6500. Example 4-21 illustrates performing troubleshooting of dCEF by examining the hardware-switching adjacency table on a Catalyst 4500 switch.

Example 4-21 *Viewing a Hardware-Switching Adjacency Table*

```
Router-2#show platform hardware ip adjacency host ip 172.18.114.1
32757: src: 00:08:A3:78:BD:FF dst: 00:08:A3:7F:CB:7C
       lbc: 0 vlan: 114 port: 248 (Gi1/1) sifact: FwdToCpu
       ifaid: 4086 packets: 290226195 bytes: 249761629484085 size: 1 refs: 3
       age: 2856047.074sec umda: flood
       vlanId: 114 single: true shar: true
       ifa: (4086) int: Vl114 (vlan 114) 172.18.114.1 00:08:A3:7F:CB:7C normal
```

Troubleshooting Punt Adjacencies

In Example 4-22, suppose that the **show ip cef** command output yielded the following as the result of a configuration change, such as routing the traffic over a tunnel interface.

Example 4-22 *Determining Whether the FIB Prefix Points to Punt Adjacency*

```
Router-2#show ip cef 10.18.118.1
10.18.118.0/24, version 31, epoch 0
0 packets, 0 bytes
  via 192.168.1.1, 0 dependencies, recursive
    next hop 192.168.1.1, Tunnel1 via 192.168.1.0/24
    valid punt adjacency
```

The CEF table output indicates that the entry is a valid punt adjacency. The term *punt* in Cisco IOS refers to sending a packet to the next-level switching process. By default, Cisco IOS always uses the fastest switching method possible; however, the fastest switching method usually does not support advanced features such as Network Address Translation (NAT) and policy-based routing (PBR) early in the product life cycle. Therefore, to handle these cases effectively, Cisco IOS punts the packet to the next-level switch method that can switch the frame. The following list illustrates the typical switching method hierarchy, with the fastest, most effective switching method listed first:

- Hardware-based dCEF (hardware-based switching on line cards)
- Hardware-based CEF (hardware-based, centralized forwarding typically found on Cisco Catalyst switches)
- PXF switching
- Software-based CEF
- Software-based fast switching
- Software-based process switching

Generally, the fastest and most effective switching methods generally lag software-based CEF in feature support. Therefore, when using new and unique features, Cisco IOS generally supports the feature in software first in low- to mid-range routers and switches. High-end routers and switches only support advanced features in hardware because the throughput required for the high-end routers and switches exceeds the capability of software-based switching methods. Use the **show cef not-cef-switched** command to view packets that are not CEF switched. In later code, the **show ip cef switching statistics** command gives detailed information of why a pass occurs and replaces the **show cef not** command. Example 4-23 illustrates an example of the **show cef not-cef-switched** command followed by Table 4-4, which illustrates descriptions for each of fields.

Example 4-23 *Sample Output from the* **show cef not-cef-switched** *Command*

```
Router-2#show cef not-cef-switched
CEF Packets passed on to next switching layer
Slot  No_adj No_encap Unsupp'ted Redirect  Receive  Options   Access     Frag
RP         4       0           0       34     2920        0        0        0
```

Table 4-4 **show cef not-cef-switched** *Command Field Descriptions*

Field	Description
Slot	Refers to the slot for the respective ingress packet counts. For Cisco IOS routers that do not support dCEF, this value is always RP.
No_adj	Indicates the number of packets dropped because of incomplete adjacency.
No_encap	Indicates the number of packets sent to the processor for ARP resolution.
Unsupp'ted	Indicates the number of packets fast-dropped by CEF (drop adjacency) because of an unsupported feature.
Redirect	Indicates the number of packets requiring ICMP redirect by process switching.
Receive	Indicates the number of packets ultimately destined to the router, or packets destined to a tunnel endpoint on the router. If the decapsulated tunnel packet is IP, the packet is CEF switched. Otherwise, packets are process switched.
Options	Indicates the number of packets with options. Packets with IP options are process switched.
Access	Indicates the number of packets punted because of an access list failure.
Frag	Indicates the number of packets punted because of fragmentation failure.
MTU	Indicates the number of packets punted because of maximum transmission unit (MTU) failure. Note: This field is not supported for IPv4 packets.

As hardware-switching components evolve, more features are being supported by CEF and dCEF. Unfortunately, most legacy software features are not supported by CEF or dCEF. Nevertheless, when troubleshooting CEF, you should note whether a software feature such

as NAT, PBR, or accounting supports CEF. To verify such features, check the platform-specific release notes, data sheets, and configuration guides.

Understanding CEF Error Messages

With current-generation Cisco IOS routers and Catalyst switches, some error messages pertaining to CEF are platform specific. As a result, you need to understand platform differences when troubleshooting CEF error messages. Nevertheless, Table 4-5 lists the general CEF EXEC error messages applicable to all Cisco IOS routers and Catalyst switches when configuring CEF and CEF features.

Table 4-5 *General CEF EXEC Error Messages*

Error Message	Troubleshooting Action
Must enable IP routing first.	When enabling CEF, IP routing must be enabled first.
Cannot disable CEF on this platform.	When disabling CEF, some platforms, such as the Catalyst 6500, do not support disabling CEF.
CEF is not enabled.	When attempting to configure a global configuration command that requires CEF, CEF must first be enabled.
CEF not enabled on this interface. Enable first.	When attempting an interface configuration that requires CEF, CEF must be enabled first.

A useful tool when working with CEF error messages is the Error Message Decoder on Cisco.com found at the following website (you must be a registered Cisco.com user):

www.cisco.com/cgi-bin/Support/Errordecoder/index.cgi

Troubleshooting Commands Reference

Table 4-6 reviews the most important commands discussed and illustrated in this chapter.

Table 4-6 *Important CEF Troubleshooting Commands*

Command	Description
debug ip cef	Enables debugging with IP CEF from the Cisco IOS CLI. Because this debug is intrusive, the recommended procedure is to use this debug option with optional parameters including ACLs to limit output. Furthermore, you should use this **debug** command under the supervision of a Cisco TAC engineer.
show adjacency detail	Displays the IP CEF Layer 2 rewrite information and statistical information from the perspective of the software-switching path.

continues

Table 4-6 *Important CEF Troubleshooting Commands (Continued)*

Command	Description
show arp	Displays the ARP table contents. This command is the first step in troubleshooting adjacency issues because the adjacency table is built from the ARP table.
show ip cef	Displays the IP CEF information from the software-switching path perspective. This command displays the prefix, next hop, and next-hop interface (outgoing interface) for verification. This command does not clearly illustrate the information contained through the hardware-switching path.
show ip route	Displays the IP routing table. This command is the first step in troubleshooting IP routing or IP CEF issues because the CEF tables are built from the IP routing table from a high-level perspective.

The following checklist reviews the important points of troubleshooting a software-based CEF issue:

- Never assume that you are experiencing a Cisco IOS CEF issue; always troubleshoot from Layer 1 (physical layer) up.
- In terms of troubleshooting CEF, verify the IP routing table and ARP table as a first step, because the CEF tables are built from those tables.
- In Cisco IOS, use the **show interface** command to verify physical layer connectivity as the initial step in troubleshooting any IP connectivity issue.
- Using the ICMP echo command **ping** from the Cisco IOS CLI always initiates echoes using the software-switching path.
- After verifying the ARP and IP routing table in Cisco IOS, verify the software-switching path for CEF by verifying the rewrite and next-hop information from the **show ip cef detail** and **show adjacency** commands.
- In controlled environments on select platforms that do not support dCEF, disabling CEF is an available troubleshooting option.
- Most mid- to high-end routers and all Catalyst switches do not support disabling CEF because CEF is required for normal operation.
- To determine why CEF is dropping ingress packets, use the **show cef drop** command.
- When troubleshooting CEF punt adjacencies, use the **show cef not-cef-switched** command to determine which CEF drop category the drops fall into.

Summary

This chapter covered many topics outside of CEF and presented a model that eventually leads to CEF troubleshooting. CEF is a scapegoat for many IP connectivity issues. The most common IP connectivity issues that are not CEF issues occur at the physical layer and Layer 2, such as spanning-tree or HSRP failures. Therefore, you should troubleshoot based on the OSI model. If a device uses hardware switching, conduct additional troubleshooting to isolate the issue.

References

Cisco, CEF Quick Help Guide, www.cisco.com/warp/customer/732/Tech/switching/docs/cef_troubleshooting.pdf. (This link requires Cisco.com registration.)

Cisco, Configuring CEF, www.cisco.com/en/US/customer/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800ca7cc.html. (This link requires Cisco.com registration.)

Cisco, Configuring Cisco Express Forwarding, www.cisco.com/en/US/products/sw/iosswrel/ps1831/products_configuration_guide_chapter09186a00800ca6ca.html.

Cisco, How to Verify Cisco Express Forwarding Switching, www.cisco.com/en/US/customer/products/sw/iosswrel/ps1828/products_tech_note09186a00801e1e46.shtml. (This link requires Cisco.com registration.)

Cisco, Troubleshooting Incomplete Adjacencies with CEF, Document ID: 17812, www.cisco.com/warp/customer/105/cef_incomp.html. (This link requires Cisco.com registration.)

Cisco, Troubleshooting Cisco Express Forwarding, Document ID: 26083, www.cisco.com/warp/customer/105/trouble_cef.html. (This link requires Cisco.com registration.)

Cisco, Troubleshooting Prefix Inconsistencies with Cisco Express Forwarding, www.cisco.com/en/US/customer/tech/tk827/tk831/technologies_tech_note09186a00800946f7.shtml. (This link requires Cisco.com registration.)

Cisco, Cisco Express Forwarding: Command Changes, www.cisco.com/en/US/products/sw/iosswrel/ps1838/products_feature_guide09186a008029b100.html.