

Troubleshooting Incomplete Adjacencies with CEF

Document ID: 17812

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

What Is an Adjacency?

Types of Adjacency

Adjacency Discovery

Reasons for Incomplete Adjacencies

- No ARP Entry
- Not Deleted After Marked Incomplete

Known Issues

Related Information

Introduction

Network nodes in the network are considered adjacent if they can reach each other with a single hop across a link layer. This document provides tips on how to troubleshoot incomplete adjacencies, as the output of the **show ip cef adjacency** command shows when Cisco Express Forwarding (CEF) is enabled on an interface.

```
Router#show ip cef adjacency serial 4/0/1 10.10.78.69 detail
IP Distributed CEF with switching (Table Version 2707655)
 130703 routes, 0 reresolve, 0 unresolved (0 old, 0 new), peak 39517
 130703 leaves, 9081 nodes, 26227536 bytes, 2685255 inserts, 2554552 invalidations
 949 load sharing elements, 318864 bytes, 71787 references
 universal per-destination load sharing algorithm, id 9E3B1A95
 2 CEF resets, 23810 revisions of existing leaves
 Resolution Timer: Exponential (currently 1s, peak 16s)
 22322 in-place/0 aborted modifications
 refcounts: 2175265 leaf, 1972988 node

Table epoch: 0 (17 entries at this epoch)

Adjacency Table has 112 adjacencies
 4 IPv4 incomplete adjacencies
```

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Cisco Express Forwarding (CEF)
- Configuring Cisco Express Forwarding

- How to Verify Cisco Express Forwarding Switching

Components Used

The information in this document is based on the Cisco IOS® Software Release 12.3(3).


The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

What Is an Adjacency?

CEF describes a very high speed switching mechanism that a router uses to forward packets from the inbound to the outbound interface. CEF uses two sets of data structures or tables, which it stores in router memory:

- **Forwarding Information Base (FIB)**  Taken from the common International Organization for Standardization (ISO) usage, an FIB describes a database of information used to make forwarding decisions. It is conceptually similar to a routing table or route-cache, although it is very different from a routing table in implementation.
- **Adjacency table** Two nodes in the network are considered adjacent if they can reach each other using a single hop across a link layer. For example, when a packet arrives at one of the router's interfaces, the router strips off the data-link layer framing and passes the enclosed packet to the network layer. At the network layer, the destination address of the packet is examined. If the destination address is not an address of the router's interface or the all hosts broadcast address, then the packet must be routed.

At a minimum, each route entry in the database must contain two items:

- ◆ **Destination address** This is the address of the network the router can reach. The router may have more than one route to the same address.
- ◆ **Pointer to the destination** This pointer indicates that the destination network is directly connected to the router, or it indicates the address of another router on a directly-connected network towards the destination. That router, which is one hop closer to the destination, is the next-hop router. An adjacency represents the pointer to the destination.

This example uses an Ethernet interface of a router (for example R1) configured with an IP address of 172.16.81.98 and a simple default static route that points all destinations to the Ethernet interface of a neighboring router R2, with an IP address of 172.16.81.1 as the next hop. In general, CEF needs to be enabled on the incoming interface for packets to be CEF switched. Since CEF makes the forwarding decision on input, use the **no ip route-cache cef** command on the ingress interface to disable CEF.

Note: In fast-switching, Cisco IOS builds a fast-switching cache entry after it switches a packet. For example, a packet that comes on a process-switched interface and is sent out through a fast-switched interface is fast switched. Issue the **no ip route-cache** command on the egress interface to disable fast switching. This is in contrast to CEF.

1. Use the **show ip route** command to view the contents of the IP routing table.

```
R1#show ip route
Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```

E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route
Gateway of last resort is 172.16.81.1 to network 0.0.0.0

```

```

172.16.0.0/24 is subnetted, 1 subnets
C    172.16.81.0 is directly connected, Ethernet0/0
S*  0.0.0.0/0 [1/0] via 172.16.81.1

```

```

!--- A simple default static route points all destinations to
!--- a next-hop address of 172.16.81.1.

```

2. Use the **show ip arp** or the **show arp** command to display the Address Resolution Protocol (ARP) table.

Note: The "Hardware Addr" field in the ARP table displays entries for the local interface and the next-hop interface.

```

R1#show ip arp
Protocol Address      Age (min)  Hardware Addr  Type   Interface
Internet 172.16.81.98      -         0030.71d3.1000 ARPA   Ethernet0/0
Internet 172.16.81.1       0         0060.471e.91d8 ARPA   Ethernet0/0

```

3. Use the **show adjacency ethernet 0/0 detail** and the **show adjacency ethernet 0/0 internal** commands to view the contents of the adjacency table entry.

```

R1#show adjacency ethernet 0/0 detail
Protocol Interface      Address
IP       Ethernet0/0     172.16.81.1(7)
                                0 packets, 0 bytes
                                0060471E91D8003071D310000800
                                ARP          03:57:08
                                Epoch: 1

```

```

R1#show adjacency ethernet 0/0 internal
Protocol Interface      Address
IP       Ethernet0/0     172.16.81.1(7)
                                0 packets, 0 bytes
                                0060471E91D8003071D310000800
                                ARP          03:57:00
                                Epoch: 1
                                Fast adjacency enabled
                                IP redirect enabled
                                IP mtu 1500 (0x48000082)
                                Fixup disabled
                                Adjacency pointer 0x62515AC0, refCount 7
                                Connection Id 0x0
                                Bucket 236

```

This output illustrates that in CEF, an adjacency refers to a control structure that holds Layer 2 information for an IP address on a particular interface. It contains the rewrite string that varies with the encapsulation protocol of the outbound interface. An adjacency is CEF's equivalent of an ARP entry.

This table describes key fields in the **show adjacency [interface-type interface-number] internal** command.

Field	Description
172.16.81.1(7)	IP address of the next-hop interface. The value in parenthesis refers to the "refCount" or the number

	of times that this adjacency is pointed to by FIB entries. The same value appears later in the entry.
0 packets, 0 bytes	Use the ip cef accounting command to enable packet and byte counters.
0060471E91D8003071D310000800	The first twelve characters are the MAC address of the destination next-hop interface. The next twelve characters represent the MAC address of the source interface of the packet. (In other words, the outbound interface of the local router). The last four characters represent the well-known Ethertype value 0x0800 for IP (with Advanced Research Projects Agency (ARPA) encapsulation).
003071D310000800	MAC address and well-known Ethertype value 0x0800 for IP (with ARPA encapsulation) of the source interface of the packet. (In other words, the outbound interface of the local router).
ARP 03:57:00	ARP indicates how the entry is discovered. The timestamp indicates how long to go before the entry times out.
Epoch: 1	CEF adjacency table Epoch information. Use the show ip cef epoch command to display the epoch information for the adjacency table and all FIB tables.
Fast adjacency enabled	An FIB entry caches an adjacency for a next-hop interface when not doing load-sharing over multiple active paths. A fast adjacency facilitates faster switching of packets.

Adjacency pointer 0x62515AC0	
refCount 7	The number of references to the adjacency that are currently stored in the router's memory. There is one for each corresponding entry in the CEF table, plus a few others for a variety of reasons (such as one for the code that performs the show adjacency command).
Connection Id 0x0	
Bucket 236	

Types of Adjacency

Adjacency Type	Adjacency Processing
Null adjacency	Packets destined for a Null0 interface are dropped. This can be used as an effective form of access filtering.
Glean adjacency	When a router is connected directly to several hosts, the FIB table on the router maintains a prefix for the subnet rather than for the individual host prefixes. 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.
Punt adjacency	Features that require special handling or features that are not yet supported in conjunction with CEF switching paths are forwarded to the next switching layer for handling. Features that are not supported are forwarded to the next higher switching level.
Discard adjacency	Packets are discarded.
Drop adjacency	Packets are dropped, but the prefix is checked.
Cached Adjacency	Cached Adjacency is the Acknowledgement update received for the adjacency packet sent.

Adjacency Discovery

Adjacencies are added to the table either through indirect manual configuration or dynamically, when discovered through a mechanism like ARP or using a routing protocol, such as BGP and OSPF, which forms neighbor relationships. If an adjacency is created by the FIB and is not discovered dynamically, then the Layer 2 addressing information is not known and the adjacency is considered incomplete. Once the Layer 2 information is known, the packet is forwarded to the route processor, and the adjacency is determined through

ARP.

ATM and Frame Relay interfaces can be configured as point-to-point or as a multipoint. The number of the type of adjacencies varies with the configuration:

- **Point-to-point interface** Uses a single adjacency for the interface.
- **Multipoint interface** Uses a unique adjacency or Layer 2 rewrite structure for each host IP address. The information to complete the adjacency comes from IP ARP, static ATM, or Frame Relay map statements, and inverse ARP on ATM and Frame Relay.

```
Router#show adjacency serial 0 detail
Protocol Interface      Address
IP         Serial0              140.108.1.1(25)
                                0 packets, 0 bytes
                                18410800
                                FR-MAP      never
                                Epoch: 1
IP         Serial0              140.108.1.2(5)
                                0 packets, 0 bytes
                                18510800
                                FR-MAP      never
                                Epoch: 1
```

When an ATM interface supports more than one permanent virtual circuit (PVC) on an interface, the "incomplete" error indication can appear for up to one minute, but it should not persist.

Note: In addition to regular adjacencies, CEF also supports five adjacency types that require special handling. These types are described in the Adjacency Types That Require Special Handling section of the Cisco Express Forwarding Overview and are outside the scope of this document.

Reasons for Incomplete Adjacencies

There are two known reasons for an incomplete adjacency:

- The router cannot use ARP successfully for the next-hop interface.
- After a **clear ip arp** or a **clear adjacency** command, the router marks the adjacency as incomplete. Then it fails to clear the entry.
- In an MPLS environment, IP CEF should be enabled for Label Switching. Interface level command **ip route-cache cef**

The symptoms of an incomplete adjacency include random packet drops during a ping test. Output drops result from throttling the rate at which CEF punts the arriving packets to the CPU. Use the **debug ip cef** command to view CEF drops due to an incomplete adjacency.

```
Router#
*Oct 11 17:08:03.275: CEF-Drop:
Stalled adjacency for 192.168.10.2 on Serial0/1/3 for
destination 192.168.11.1
*Oct 11 17:08:03.275: CEF-Drop:
Packet for 192.168.11.1 -- encapsulation
*Oct 11 17:08:05.307: CEF-Drop:
Stalled adjacency for 192.168.10.2 on Serial0/1/3 for
destination 192.168.11.1
*Oct 11 17:08:05.307: CEF-Drop:
Packet for 192.168.11.1 -- encapsulation
```

In addition, use the **show cef drop** command several times and look for an incrementing value for the 'Encap_fail' counter. Refer to the **show cef** commands for more information.

No ARP Entry

When CEF cannot locate a valid adjacency for a destination prefix, it punts the packets to the CPU for ARP resolution and, in turn, for completion of the adjacency. In rare cases, the adjacency persists in an incomplete state. For example, if the ARP table already lists a particular host, then punting it to the process level does not trigger an ARP.

Determine whether an ARP entry exists in order to troubleshoot this problem. Use these commands and specify a specific IP address:

- **show arp** or **show ip arp**
- **show adjacency**

Use the **debug arp** command to confirm that the router sends an ARP request.

```
Router#ping 10.12.241.4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.12.241.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
Router#
.Aug 21 18:59:07.175 PDT:
IP ARP:
creating incomplete entry for IP address:10.12.241.4 interface FastEthernet0/1
.Aug 21 18:59:07.177 PDT: IP ARP: sent req src 10.12.241.252 0006.529c.9801,
dst 10.12.241.4 0000.0000.0000 FastEthernet0/1
.Aug 21 18:59:07.180 PDT: IP ARP throttled out the ARP Request for 10.12.241.4
.Aug 21 18:59:09.182 PDT: IP ARP: sent req src 10.12.241.252 0006.529c.9801,
dst 10.12.241.4 0000.0000.0000 FastEthernet0/1
.Aug 21 18:59:09.183 PDT:
IP ARP throttled out the ARP Request for 10.12.241.4
```

When the ping process tries to send the first packet and does not see an ARP entry, it initiates an ARP request. It continues to try to send the packet, and then drops the packet after a defined wait period. When an ARP response is received and the ARP entry is completed using a background process, the ping success rate is 100 percent.

Not Deleted After Marked Incomplete

When adjacency information needs to be changed, the adjacency aging logic removes an entry in two stages:

- First it changes the status of the entry from complete to incomplete.

```
Router#show adjacency
Protocol Interface Address
IP Serial0 10.10.10.2(2) (incomplete)
IP Serial0 10.10.10.3(7)
IP Ethernet0 172.16.81.1(7)
```

- Then, at the next one-minute interval, the adjacency walker process "wakes up" and completes the deletion.

```
Router#show adjacency
Protocol Interface Address
IP Serial0 10.10.10.3(7)
IP Ethernet0 172.16.81.1(7)
```

In distributed CEF mode, the process on the RP informs the line cards to complete the deletion. This sequence illustrates that a window of up to 60 seconds exists in order for a transient incomplete

adjacency to exist.

Known Issues

On a Frame Relay interface, configuring a static map statement prompts CEF to add a host prefix entry to the CEF table. Originally, CEF did not consider whether the PVC was in an "ACTIVE" status before creating the entry. This issue is resolved in Cisco bug ID CSCdr71258 (registered customers only) .

In addition, after attaching to and then removing an interface from a Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) route forwarding (VRF) instance, CEF sets the adjacency to incomplete. However, the Frame Relay dynamic map entry is not cleared. When the IP address is reapplied, the dynamic mapping still exists. This prevents the adjacency from ever being completed. Issue the **clear frame-relay-inarp** command when the IP address is removed (for example when the VRF is applied) to avoid this problem. The IP address can then be reapplied, and the adjacency is completed as soon as the dynamic map is recreated.

Related Information

- [How to Verify Cisco Express Forwarding Switching](#)
- [Configuring Cisco Express Forwarding](#)
- [Cisco Express Forwarding Overview](#)
- [Cisco Express Forwarding \(CEF\) Technology Support Page](#)
- [IP Switching Technology Support Page](#)
- [Technical Support & Documentation – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jun 07, 2006

Document ID: 17812
