

Troubleshooting Prefix Inconsistencies with Cisco Express Forwarding

Document ID: 14540

Contents

Introduction

Prerequisites

Requirements

Components Used

Conventions

CEF Forwarding Tables

What is an Inconsistency?

Inconsistency Checker Overview

Troubleshooting CEF Inconsistencies

Troubleshooting Commands

How to Reset the Consistency Checker

False Inconsistencies

Known Issues

Related Information

Introduction

This document describes the consistency checker now available for the Cisco 7500 and 12000 series routers running the distributed form of Cisco Express Forwarding (CEF). The consistency checker, introduced in Cisco IOS® Software Release 12.0(15)S and other release trains, is designed to detect when forwarding information on line cards and the route processor (RP) lose synchronization. Cisco IOS reports the following log messages when the checker detects a problem:

```
%FIB-4-RPPREFIXINCONST2: RP missing prefix for  
133.160.0.0/16 (present in routing table)
```

```
%FIB-4-RPPREFIXINCONST2: RP missing prefix for  
133.160.0.0/16 (present in routing table)
```

```
%FIB-4-LCPREFIXINCONST2: Slot 1 missing prefix entry for 64.0.17.0/32
```

This document also provides troubleshooting tips on CEF inconsistencies.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on these software and hardware versions:

- Cisco IOS Software Release 12.0(15)S and later

- Cisco 7500 and 12000 series routers

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

For more information on document conventions, refer to the Cisco Technical Tips Conventions.

CEF Forwarding Tables

To understand what the consistency checker is reviewing, you first need to understand and define the CEF forwarding tables.

CEF describes a 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)** – Describes a database of information used to make forwarding decisions. It is conceptually similar to a routing table or route-cache, although its implementation is different.
- **Adjacency** – Two nodes in the network are said to be adjacent if they can reach each other via a single hop across a link layer.

The FIB table is updated when one of the following occurs:

- The Address Resolution Protocol (ARP) cache entry for the next hop changes, is removed, or times out.
- The routing table entry for the prefix changes or is removed.
- The routing table entry for the next hop changes or is removed.

The Cisco 7500 and 12000 series routers support distributed CEF (dCEF), in which the line cards make the packet forwarding decisions using locally stored copies of the same FIB and adjacency tables as the RP. The tables between the RP and the line cards must remain synchronized. Any changes to the RP's tables must be forwarded to the line cards.

Inter-Process Communication (IPC) is the protocol used by routers that support distributed packet forwarding. CEF updates are encoded as eXternal Data Representation (XDR) information elements inside IPC messages. The following diagram illustrates the CEF data structure distribution mechanism.



What is an Inconsistency?

There are two types of inconsistencies:

- Missing information, such as a particular prefix, on a line card.
- Different information, such as different next hop IP addresses, on the line card.

```

router#show ip cef 24.20.84.32
24.16.0.0/13, version 833173, cached adjacency to POS6/0
0 packets, 0 bytes
Flow: AS 6172, mask 13
via 4.24.234.153, 0 dependencies, recursive
next hop 4.24.234.153, POS6/0 via 4.24.234.152/30
valid cached adjacency
router#execute-on all show ip cef 24.20.84.32
===== Line Card (Slot 1) =====
24.16.0.0/13, version 408935, cached adjacency 0.0.0.0
0 packets, 0 bytes
Flow: AS 6172, mask 13
via 157.130.213.1, 0 dependencies, recursive
next hop 157.130.213.1, POS1/0.500 via 157.130.213.0/30
valid cached adjacency
===== Line Card (Slot 2) =====
24.16.0.0/13, version 13719, cached adjacency 0.0.0.0
0 packets, 0 bytes
Flow: AS 6172, mask 13
via 157.130.213.1, 0 dependencies, recursive
next hop 157.130.213.1, POS1/0.500 via 157.130.213.0/30
valid cached adjacency

```

One indication of a CEF inconsistency is a difference in the "CEF-ver" column of the **show cef linecard** command, as shown in the following output.

```

7505-2A#show cef linecard
CEF table version 35, 11 routes
Slot CEF-ver   MsgSent   XDRSent   Window   LowQ   MedQ   HighQ   Flags
1      0           0           0 LC wait   0      0      0 disabled
2      31          95          422      24     0      0      0 up, sync
3      34          105         544      24     0      0      0 up, sync

```

The output of the **show ip cef summary** command on the RP and LC should show the same number of routes and adjacencies when the following conditions are true:

- All line cards are in the "up, sync" state.
- XDR queues on both the RP and LC are empty ("LowQ/MedQ/HighQ" columns in **show cef linecard** on the RP and the "RP messages to be processed" data in **show cef linecard** on the linecard).

Note: An exception is Cisco 12000 Series' Engine 2 LCs, where the packet switching ASIC (PSA) installs additional routes to implement ACLs.

Cisco IOS Software Release 12.0(22)S contains the CEF consistency checkers v2 (present in Cisco IOS Software Release 12.1E), which lets you execute the command **show ip cef inconsistency now** to report any problems.

Inconsistency Checker Overview

As updates occur to the routing databases on the RP and line cards, inconsistencies may result due to the asynchronous nature of the distribution mechanism for these databases. CEF supports passive and active consistency checkers that run independently. The following table describes these checkers.

Detection Mechanism	Operates	Description
Lc-detect	On Line Card	Operates on the line card by retrieving IP prefixes found missing from its FIB table. If IP prefixes are missing, the line card cannot forward packets for these addresses. Lc-detect then sends IP prefixes to the RP for confirmation. If the RP detects that it has the relevant entry, an inconsistency is detected and an error message is displayed. Also, the RP sends a signal back to the line card confirming that the IP prefix is an inconsistency.
Scan-lc	Line Card	Operates on the line card by looking through the FIB table for a configured time period and sending the next n prefixes to the RP. The RP does an exact lookup. If it finds the prefix missing, the RP reports an inconsistency. Finally, the RP sends a confirmation back to the line card.
Scan-rp	Route Processor	Operates on the RP (opposite of the scan-lc) by looking through the FIB table for a configured time period

		and sending the next n prefixes to the line card. The line card does an exact lookup. If it finds the prefix missing, the line card reports an inconsistency and finally signals the RP for confirmation.
Scan-rib	Route Processor	Operates on all RPs (even non distributed), and scans the RIB to ensure that prefix entries are present in the RP FIB table.

Use the following commands to configure the enabled consistency checkers and related variables.

- **ip cef table consistency-check type [period] [count]** – Controls general parameters for the checkers.
- **ip cef table consistency-check** – Enables or disables supported types and controls the period of scans and prefixes scanned (not for lc-detect). The consistency checker is disabled by default.

Troubleshooting CEF Inconsistencies

Inconsistencies should never happen, and any inconsistencies should be investigated. Use the following CEF **debug** and **show** commands when troubleshooting.

Certain **show** commands are supported by the Output Interpreter Tool (registered customers only) , which allows you to view an analysis of **show** command output.

Note: Before issuing **debug** commands, please see Important Information on Debug Commands.

- **show ip cef inconsistency records detail** – Displays statistics on the detection mechanisms. Also records detailed information for a number (currently 4) of confirmed, but different, inconsistencies.
- **show ip cef inconsistency** – Displays a summary of the state.

```
Table consistency checkers (settle time 15s)
lc-detect: running
0/0/0 queries sent/ignored/received
scan-lc: running [100 prefixes checked every 60s]
0/0/1053 queries sent/ignored/received
scan-rp: running [100 prefixes checked every 60s]
1053/0/0 queries sent/ignored/received
scan-rib: running [1000 prefixes checked every 60s]
81/0/81 queries sent/ignored/received
Inconsistencies: 0 confirmed, 0/4 recorded
```

- **show ip cef inconsistency records** – Use the records keyword to dump any recorded inconsistencies.
- **show ip cef inconsistency records detail and execute-on slot**

```
router#exec slot 2 show ip cef inconsistency records detail
===== Line Card (Slot 2) =====
```

```
Table consistency checkers (settle time 65s)
lc-detect: running
0/0/0 queries sent/ignored/received
scan-lc: running [100 prefixes checked every 60s]
1289156/0/0 queries sent/ignored/received
scan-rp: running [100 prefixes checked every 60s]
0/0/1068308 queries sent/ignored/received
scan-rib: running [1000 prefixes checked every 60s]
0/0/0 queries sent/ignored/received
```

```

Inconsistencies: 340 confirmed, 1/4 recorded
Test table insert mode: normal
Test table remove mode: normal
----- Inconsistency record 0 -----
Prefix entry for 192.168.3.10/32 present on RP, missing on slot 2
Detected at 00:01:46.736 by scan-rp
Event log entries relevant for 192.168.3.10:
+00:00:00.000: *.*.*./*          New FIB table          [OK]
    0x403FA4E8 0x403FA530 0x4009C1FC 0x4009C1E8
+00:00:03.092: *.*.*./*          Flush ADJ table        [OK]
    0x404000B0 0x4040EEC0 0x4040F100 0x40401F64 0x404021AC 0x4040229C
0x404029C8 0x4009C1FC 0x4009C1E8
+00:00:03.100: *.*.*./*          Flush FIB table        [OK]
    0x404039D0 0x40401F4C 0x404021AC 0x4040229C 0x404029C8 0x4009C1FC
0x4009C1E8
+00:00:03.124: *.*.*./*          New FIB table          [OK]
    0x404039D8 0x40401F4C 0x404021AC 0x4040229C 0x404029C8 0x4009C1FC
0x4009C1E8
First event occurred at 00:00:07.600 (2w5d ago)
Last event occurred at 00:00:10.724 (2w5d ago)

```

Note: There is no SNMP access to the inconsistency records. This feature may be added to an upcoming release of Cisco IOS.

Troubleshooting Commands

- **clear cef linecard** – If an IP prefix is missing from a line card, use this command to re-install the prefix in the line card FIB.
- **clear ip route** – If a prefix is missing from the RP, use this command to re-install the prefix in the RP FIB.

When the consistency checker finds a problem, the following command outputs are needed to accurately troubleshoot the issue.

- **show ip cef** – Displays entries in the FIB on the RP.
- **exec all show ip cef** – Displays CEF FIB values on line cards.
- **show tech cef** – Provides information for technical support on CEF.
- **show ip cef inconsistency records detail** – Displays CEF FIB inconsistency details on the RP.
- **exec slot show ip cef inconsistency records detail** – Displays CEF FIB inconsistency details on the line cards.
- **no ip cef table consistency-check** – Turns off the checkers.
- **debug ip cef table consistency-checkers** – Causes query and check events to be debugged.

How to Reset the Consistency Checker

To clear a CEF inconsistency, use the **clear ip cef inconsistency** command. To turn off the consistency checker, use the **no ip cef table consistency-check** command. It is important to note that turning off consistency checker does not fix the problems reported. The system continues to run with inconsistencies, potentially leading to unexpected behavior.

False Inconsistencies

In rare conditions, the original version of the CEF consistency checker may report a false positive. This problem is caused by temporary timing windows when the CEF databases are being updated (particularly during very large updates) and new information is being distributed from the RP to the line cards. Although false positive messages are heavily rate limited – and thus are more of a nuisance than an impact to the CPU – Cisco recommends disabling the CEF consistency checker unless troubleshooting a specific consistency

problem.

To minimize the chances of false positives for prefixes in the process of being updated, you can tune the settle-time, which is the delay imposed on the detection. Use the **ip cef table consistency-check [settle-time]** command. This command relies on an event log being present; otherwise, the settle-time is effectively 0.

A default value of 65 seconds was selected to prevent phantom detection of adjacency prefixes (/32s for ARP entries) missing on the RP. An ARP delete happens in two stages on the RP:

1. The adjacency is marked as incomplete and the adjacency prefix is deleted.
2. The one minute adjacency walker deletes the adjacency and informs the line cards to do the same.

This process leads to a window of up to 60 seconds when the adjacency prefix is present on the line cards, but missing on the RP.

Note: In Version 2.0 of the consistency checker, the issue of false inconsistencies has been fixed.

Known Issues

The following lists known Cisco bugs with the CEF consistency checker. This list is not meant to be exhaustive.

- Cisco bug ID CSCdt18447 (registered customers only) CEF event logger/consistency checkers hog CPU
- Cisco bug ID CSCuk23390 (registered customers only) lc-detect causes malloc fail at interrupt level
- Cisco bug ID CSCuk23290 (registered customers only) CEF inconsistency event recorder can misallocate storage
- Cisco bug ID CSCdt04645 (registered customers only) Resolves a problem with a memory leak (which affects anything that can run distributed CEF, including the lc-stat consistency checker process).

Related Information

- [Troubleshooting Cisco Express Forwarding-Related Error Messages](#)
- [Cisco Express Forwarding Technical Support Pages](#)
- [Technical Support – Cisco Systems](#)

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

© 2013 – 2014 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Aug 10, 2005

Document ID: 14540
