

Table of Contents

<u>Troubleshooting Fabric Ping Timeouts and Failures on the Cisco 12000 Series Internet Router</u>	1
<u>Document ID: 12421</u>	1
<u>Introduction</u>	1
<u>Before You Begin</u>	1
<u>Conventions</u>	1
<u>Prerequisites</u>	1
<u>Components Used</u>	1
<u>Background</u>	2
<u>Symptoms</u>	2
<u>Troubleshooting</u>	3
<u>Useful debug and show Commands</u>	6
<u>Information to Collect if you Open a TAC Case</u>	8
<u>Related Information</u>	9

Troubleshooting Fabric Ping Timeouts and Failures on the Cisco 12000 Series Internet Router

Document ID: 12421

- Introduction**
- Before You Begin**
 - Conventions
 - Prerequisites
 - Components Used
 - Background
- Symptoms**
- Troubleshooting**
- Useful debug and show Commands**
- Information to Collect if you Open a TAC Case**
- Related Information**

Introduction

This document explains how to troubleshoot fabric ping timeouts and failures on the Cisco 12000 Series Internet Router. Such failures are indicated by the following error messages:

```
%GRP-3-FABRIC_UNI: Unicast send timed out (3)
```

and

```
%GRP-3-COREDUMP: Core dump incident on slot 3, error: Fabric ping failure (seq:29192)
```

Before You Begin

Conventions

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

Prerequisites

There are no specific prerequisites for this document.

Components Used

The information in this document is based on the hardware version below.

- Cisco 12000 Series Internet Routers

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Background

The GRP and line cards (LCs) in the Cisco 12000 Series Internet Router connect through a crossbar switch fabric, which provides a high-speed physical path for most inter-card communication. Among the messages passed between the GRP and the line cards over the switch fabric are included actual packets being routed and received, forwarding information, traffic statistics, and most management and control information. Thus, it is important for the GRP to ensure that this path is operating correctly.

Fabric pings are one of four applications that run between the GRP and the switch fabric. Inter-Processor Communication (IPC), network packets, and code downloads are the others. Fabric pings are implemented to provide part of a failure detection algorithm and keep-alive mechanism implemented using buffers on the Maintenance Bus (MBUS) and pings through the line card fabric interfaces.

The Cisco Cell Segmentation and Reassembly (CSAR) Fabric Interface drivers on the GRP handle messages that are to be sent and received between the switch fabric and the GRP. This includes fabric pings. Fabric pings are generated by software and are sent from the primary GRP to each line card every six seconds. Every time a line card receives a ping request from the GRP, the LC sends back a reply to the GRP. If the GRP does not receive any reply to five consecutive fabric pings (30 seconds total time), it declares the line card dead and resets it through the Maintenance BUS (MBUS).

Most of the time, the line card is simply too busy to respond to the fabric ping requests from the GRP. These fabric ping failures might also be caused by a faulty fabric or a bug in the Cisco IOS® software. All possible causes of fabric ping failures are detailed in the troubleshooting section below.

Fabric ping timeouts occur when the Gigabit Route Processor (GRP) detects that a ping request is stuck in the ToFab queue (towards the switch fabric) of the Cisco Cell Segmentation and Reassembly (CSAR) application-specific integrated circuit (ASIC). This ASIC is responsible for slicing the packets into Cisco Cells before sending them through the switch fabric to the output line card (LC).

Fabric ping failures occur when either a line card or the secondary GRP fails to respond to a fabric ping request from the primary GRP over the switch fabric. Such failures are a problem symptom that should be investigated.

Symptoms

As explained in the Background section, the GRP sends a fabric ping to the line cards every six seconds, and the line cards must respond. When the GRP fails to receive a reply to five consecutive fabric pings, it resets the line card by sending a request message over the maintenance bus (MBUS), and reports a software-forced crash, as seen in the output of the **show context slot {#}** command.

From the the console logs or the **show log** command, you may get the following error messages before the fabric ping failure message:

```
%GRP-3-FABRIC_UNI: Unicast send timed out (3)
%GRP-3-FABRIC_UNI: Unicast send timed out (3)
%GRP-3-FABRIC_UNI: Unicast send timed out (3)
```

where the number (3) represents the line card slot towards which the primary GRP tried to send a fabric ping.

This message indicates that a packet is stuck in the ToFab queue of the CSAR ASIC on the primary GRP. If anything is stuck in either of the two CSAR buffers for more than 100 milliseconds (msecs), the buffer is flushed and a timeout message is generated.

If the GRP sends its fabric ping request message, but either the line card does not answer, or the line card answers but the switch fabric is faulty so it loses the message, you will not see this message prior to the fabric ping failure message. Therefore, if you get the error message "%GRP-3-FABRIC_UNI", this means that something could not be transmitted to a slot over the fabric for 100 or 200 msec. It might be that, because of the %GRP-3-FABRIC_UNI, you cannot send the keep-alives to the LC and you end up with a fabric ping failure after, in this case, 30 seconds. However, you can get fabric ping failures without the "%GRP-3-FABRIC_UNI" and vice versa.

The primary GRP may determine that a line card or secondary GRP has degraded to a point that a diagnostic core dump is appropriate. At this time, the GRP sends a message over the MBUS to the line card and asks the line card CPU to crash, so a core dump can be obtained.

```
%LCINFO-3-CRASH: Line card in slot 3 crashed
%GRP-3-COREDUMP: Core dump incident on slot 3, error: Fabric ping failure (seq:29192)
```

The line card creates the core dump if it is configured with the **exception crashinfo** and related commands (see Configuring a Core Dump on a GSR Line Card for GSR-specific information on configuring core dumps). The qualifying string in the output of the **show context slot {#}** command indicates the reload reason. In the case of a fabric ping failure, the reason is always "Software-forced crash".

```
CRASH INFO: Slot 1, Index 1, Crash at 00:42:45 KST Mon Mar 12 2001
VERSION:
GS Software (GLC1-LC-M), Version 12.0(18)ST, EARLY DEPLOYMENT RELEASE SOFTWARE (fc1)
TAC Support: http://www.cisco.com/tac
Compiled Thu 09-Aug-01 22:06 by nmasa
Card Type: 2 Ports OC3 Channelized to DS1/E1 , S/N CAT00400500
System exception: sig=23, code=0x24,

! --- SIG=23 indicates a software-forced crash.

context=0x41303B04
System restarted by a Software forced crash
STACK TRACE:
-Traceback= 400C3970 400C1F90 40815D5C 407D3144 400C7488
```

After the line card crashes, it sends an initial message to notify the primary GRP. The GRP then waits for the line card to send other information about the crash through the MBUS. The GRP should receive full information within a few milliseconds after receiving the first message from the line card. In the unlikely event that the subsequent crash information messages are not received by the GRP within a reasonable time limit (10 seconds), the GRP prints an error message and tells the rest of the GRP software that the line card has crashed.

Troubleshooting

During normal router operation, the primary GRP continually pings the line cards, and the line cards respond. Any ping failures are a symptom of another problem that should be investigated. These problems include:

- Problem with the Line Card
- Problem with the Switching Fabric
- Problem with the GRP
- Known issues with IPC
- Known issues with Cisco Express Forwarding (CEF)

Note: If the failure can be reproduced, configure **no service auto-reset** on the GRP. This command disables a reload of the line card at the next fabric ping failure, and allows you to attach to the line card using the

attach <slot#> command to capture relevant **show** commands.

Problem with the Line Card

- The most likely reason could be a Cisco IOS software bug in which a process is disabling interrupts long enough to miss five consecutive fabric pings. Try upgrading to the latest Cisco IOS software release in your train to avoid resolved issues. For upgrade assistance, see the Cisco IOS Upgrade Planner Tool (registered customers only).
- The line card may be exerting back-pressure for too long, so the scheduler is not allowing any traffic to be received from the switch fabric. This symptom suggests a problem with interface congestion. Use the following commands to confirm these symptoms:

- ◆ **show controller frfab queue** command on the line card. Look for a non-IPC free queue with little or no available buffers.
- ◆ **show controllers csar queue** command on the GRP. Look for non-zero values for "Max Length" and for "Max Length" to equal the "Length" value, as shown in the following sample output:

```
router#show controllers csar queue
1190 Free Q

Slot Length Max Length
0      0      7
1      0      2
2     70     70

! -- CSAR queue for slot 2 is building and reaching max length.

3      0      2
4      0      3
5      0      0
...
```

The CSAR queues up to 50 packets to a destination line card. After 50 packets, only fabric ping packets are queued. If the queue limit then increases to 70, the CSAR stops queuing all packets — including fabric pings — to the line card. Both the GRP and all line cards have 64k CSAR segmentation buffers in which to store messages. If these buffers are busy, the router uses a software hold queue to store the messages. It also sets a timer to ensure that fabric ping messages do not stay on this queue for too long.

- High CPU utilization on the line card – Commonly seen during recalculation of a large Cisco Express Forwarding (CEF) table after a massive routing table change or after a link flap and Border Gateway Protocol (BGP) session reset. The CPU can also be high when switching traffic in software. This happens mainly on Engine 0 line cards where most of the features are implemented in software. If this is the case, you may check the configuration of the line card and remove the features that could impact the CPU on the Engine 0 LC. High CPU utilization may also be due to a bug. Determine the CPU utilization with the **execute-on slot <slot#> show proc cpu** command or **execute-on slot <slot#> show tech** command if the previous command is not supported in the Cisco IOS software version which is running on the router. Consider upgrading to the latest Cisco IOS software release in your train to work around known issues.
- The line card is running out of Inter-Process Communication (IPC) buffers, which are used to exchange control messages between the line cards and the GRP. See the troubleshooting steps in Troubleshooting CEF-Related Error Messages. If your troubleshooting points to a problem with IPC, ensure that your Cisco 12000 Series Internet Router is running a minimum of Cisco IOS Software Release 12.0(18)S. This release has introduced a larger default size of 5000 for the IPC cache to enhance its stability and scalability.

- Hardware problem on the line card. It is important to note that less than 10% of fabric ping failures result from a hardware problem. Before contacting the Cisco TAC to request replacement hardware, please try the following steps:

1. Look for IPC timeout messages printed before the fabric ping failure. Also see the IPC section below.
2. Reseat the line card.
3. Power-cycle the router.
4. If you do not have physical access to the router, execute the **hw-module slot <slot #> reload** command to perform a manual reload of the line card.

Problem with the Switching Fabric

The heart of the Cisco 12000 Series Internet Router is the switch fabric circuitry, which provides synchronized gigabit speed interconnections for the line cards and the GRP. The switch fabric circuitry contains two types of cards:

- Clock and scheduler cards (CSCs)
- Switch fabric cards (SFCs)

If one of these cards is failing, the ping messages can no longer pass through the fabric. In this case, you should also see other messages pointing to the faulty fabric, such as the following:

```
%FABRIC-3-CRC: Switch card 18
```

Use the **show controllers fia** command to determine whether you have a bad CSC or SFC. Use the **execute-on all show controllers fia** command to capture output from all line cards. Compare the output from the GRP with the output from the line cards to determine whether a faulty switching fabric card needs to be replaced.

The following sample output points to a problem with sfc0 in slot 18. First try to reseat this card and then request a replacement if the crc16 error counter continues to increment.

```
Router#show controllers fia
Fabric configuration: Full bandwidth redundant
Master Scheduler: Slot 17
From Fabric FIA Errors
-----
redund FIFO parity 0    redund overflow 0    cell drops 1
crc32 lkup parity 0    cell parity 0    crc32 0
Switch cards present 0x001F Slots 16 17 18 19 20
Switch cards monitored 0x001F Slots 16 17 18 19 20
Slot: 16 17 18 19 20
Name: csc0 csc1 sfc0 sfc1 sfc2
-----
Los 0 0 0 0 0
state Off Off Off Off Off
crc16 0 0 4334 0 0

! --- Check the CRCs under SFC0 (slot 18)

To Fabric FIA Errors
-----
sca not pres 0    req error 0    uni FIFO overflow 0
grant parity 0    multi req 0    uni FIFO undrflow 0
cntrl parity 0    uni req 0    crc32 lkup parity 0
multi FIFO 0    empty DST req 0    handshake error 0
```

Problem with the GRP

In some reports of fabric ping failures, the router reported cyclic redundancy check (CRC) error messages before the failure. Check for CRCs on the switching fabric cards by using the **show controllers fia** command on the GRP and the **execute-on all show controllers fia** on the line cards. CRC errors on the GRP only (and not on any line card) point to a faulty GRP. First try to reseal the GRP and then request a replacement if the CRC errors continue to increment.

Known Issues With IPC

Problems with the Inter-Process Communication (IPC) software running between the GRP and the line cards have been resolved in various releases of Cisco IOS Software Release 12.0S. In this case, you should see some IPC-related error messages in the log, along with fabric ping timeout messages. Try running the latest Cisco IOS software release to work around known issues with IPC. Also see the Cisco IOS Upgrade Planner Tool (registered customers only) for assistance with selecting a release.

Known Issues with Cisco Express Forwarding (CEF)

See Troubleshooting CEF-Related Error Messages if the output of the **show log** command displays a message related to the CEF Forwarding Information Base (FIB) similar to the one below:

```
%FIB-3-FIBDISABLE: Fatal error, slot 2: IPC failure
```

Useful debug and show Commands

Use the following **debug** and **show** commands to troubleshoot fabric ping timeout/failure messages on the Cisco 12000 Series Internet Router:

- **debug fabric events** – Prints any errors detected by the GRP. This debug generates very few messages and only in an error condition.
- **debug fabric ping** – Prints any errors detected in the fabric ping process by the GRP. This debug generates very few messages and only in an error condition.

Capture the following commands for each reset line card. Replace X with the appropriate slot number.

- **execute-on slot X debug fabric events** – Prints errors detected by the line card in its ping responses. This command produces very few messages and only in an error condition.
- **exec slot X debug fabric ping** – Prints a message when the line card receives a fabric ping. This debug generates one line of output every second for every line card on which it is enabled.

After the line card crashes, capture the following commands from the GRP console:

- **show context all detail**
- **show fabric**
- **show controllers fia**
- **show controllers csar queue**
- **execute-on all show controllers fia**
- **show tech**
- **show log**

Capture also the following commands about the line card status:

- **execute-on slot <slot#> show proc CPU**
- **execute-on slot <slot#> show controller tofab queue**
- **execute-on slot <slot#> show controller tofab stat**
- **execute-on slot <slot#> show controller frfab queue**
- **execute-on slot <slot#> show controller frfab stat**
- **execute-on slot <slot#> show ipc stat**
- **execute-on slot <slot#> show ipc queue**
- **execute-on slot <slot#> show stack**
- **execute-on slot <slot#> show tech**

If you still encounter problems after following all the troubleshooting steps, gather all the required information above and call your Cisco TAC representative to troubleshoot further.

Here is the output from some useful **show** commands:

```
router#show controllers csar
From Fabric Error Stats
-----
0 out of order, 0 unexpected first
0 unexpected last, 0 unknown rx type, 0 corrupted pak, 0 parity
0 first/last, 0 sequence, 0 cell avail, 0 reassembly,

To Fabric Stats
-----
Slot  Tx Pkts      TX Th Pkts      Rx Pkts      Rx Th Pkts      To Fab timeout
0      580278      490214      281061      1336470      0

1      18854      66592      18390      945419      0
2      6      50824      0      896290      0
3      0      0      0      0      0
4      0      51909      0      895430      0
5      0      0      0      0      0
6      0      35113      0      880247      0
7      0      52690      0      52690      0
8      0      0      0      0      0
9      0      0      0      0      0
10     0      0      0      0      0
11     0      0      0      0      0
12     0      0      0      0      0
13     0      0      0      0      0
14     0      0      0      0      0
15     0      0      0      0      0
0 too big, 1 Buf0 free, 1 Buf1 free
0 Copy fail

Fabric access Error Stats
-----
0 parity errors, 0 bad access size, 0 invalid address
0 queue full parity, 0 flushed buffer

router#show controllers fia
Fabric configuration: Full bandwidth, nonredundant fabric
Master Scheduler: Slot 16

From Fabric FIA Errors
-----
redund fifo parity 0      redund overflow 0      cell drops 0
crc32 lkup parity 0      cell parity 0      crc32 0
Switch cards present 0x001D  Slots 16 18 19 20
```

```

Switch cards monitored 0x001D   Slots  16 18 19 20
Slot:      16      17      18      19      20
Name:      csc0      csc1      sfc0      sfc1      sfc2
-----
los       0         0         0         0         0
state    Off       Off       Off       Off       Off
crc16  0         254      0         0         0

```

! --- Check the CRC error here. In this case CSC1 in slot 17.

To Fabric FIA Errors

```

-----
sca not pres 0          req error    0          uni FIFO overflow 0
grant parity 0          multi req   0          uni FIFO undrflow 0
cntrl parity 0          uni req     0          crc32 lkup parity 0
multi FIFO  0          empty dst req 0        handshake error  0
cell parity  0

```

You can find more details on the **show controllers fia** command at How To Read the Output of the Show Controller fia Command.

```

router#show fabric
Dest      ToFab      FrFab      Bad Seq      Unexpected
Slot      Pkts      Pkts      Pkts      Pkts
-----
Slot0     26327     26327     0          0
Slot1     26325     26325     0          0
Slot2     26321     26321     0          0
Slot4     26315     26315     0          0
Slot6     26311     26311     0          0
Slot7     26334     26334     0          0

multicast timeout 0
failed pak        0

Current fabric timeout is 6000

fabric send fails 58

```

Information to Collect if you Open a TAC Case

If you still need assistance after following the troubleshooting steps above and want to create a service request with the Cisco TAC, please attach the following information to your case for troubleshooting fabric ping problems on the Cisco 12000 Series Internet Router:

- troubleshooting performed before opening the case
- **show technical-support** output (in enable mode if possible)
- **show log** output or console captures, if available
- **execute-on slot [slot #] show technical-support** for the slot that experienced the line card crash

Please attach the collected data to your case in non-zipped, plain text format (.txt). You can attach information to your case by

uploading it using the Case Query tool (registered customers only) .
If you cannot access the Case Query tool, you can attach the relevant information to your case by sending it to attach@cisco.com with your case number in the subject line of your message.

Note: Please do not manually reload or power-cycle the router before collecting the above information, if possible, as this can cause important information to be lost that is needed for determining the root cause of the problem.

Related Information

- [Configuring a Core Dump on a GSR Line Card](#)
- [Troubleshooting CEF-Related Error Messages](#)
- [How To Read the Output of the Show Controller fia Command](#)
- [Product Support – 12000 Series Internet Routers](#)
- [Technical Support – Cisco Systems](#)

All contents are Copyright © 1992–2005 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.

Updated: Jun 27, 2005

Document ID: 12421
