

# Troubleshooting Ignored Errors and No Memory Drops on the Cisco 12000 Series Internet Router

Document ID: 18003

---

## Introduction

### Prerequisites

Requirements

Components Used

Conventions

### Symptoms

Managing Outbound Interface Oversubscription

Managing Overloaded CPU on the Inbound Line Card

### Summary of Design Guidelines

### Case Study

Step 1 – Check the ToFab queues in slot 6.

Step 2 – Check the FrFab queues in slot 4.

Step 3 – Compare the output of the `show interfaces` command for the same oversubscribed interface.

Step 4 – Implement a solution.

### Cisco IOS Software Bugs

### Related Information

---

## Introduction

This document explains how to troubleshoot why the output of the `show interfaces` command on a Cisco 12000 Series Internet Router displays an increasing number of ignored errors. It also provides troubleshooting tips for an increasing number of **no mem drops** in the output of the `execute-on slot <slot#> show controllers (frfab | tofab) qm stat` command. When troubleshooting any of these errors, verify that the counter is incrementing and is not simply an historical value.

**Note:** An increasing number of input queue drops, as displayed in the `show interfaces` output, is covered separately in Troubleshooting Input Drops on the Cisco 12000 Series Internet Router.

## Prerequisites

### Requirements

This document requires an understanding of the Cisco 12000 Series Internet Router architecture, particularly the ToFab and FrFab queues. See How To Read the Output of the `show controllers frfab | tofab queue` Commands for reference.

### Components Used

The information in this document is based on the software and hardware versions below.

- Any Cisco IOS® software release which supports the Cisco 12000 Series Internet Router. Usually these are the 12.0S and 12.0ST releases.
- All the Cisco 12000 platforms are covered by this document. These include the 12008, 12012, 12016, 12404, 12410, and 12416.

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.

## Conventions

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

## Symptoms

The Cisco 12000 Series Internet Router uses a distributed architecture to ensure optimum forwarding performance. To support high forwarding rates, it maintains packet buffers on both the inbound and outbound line cards. These packet buffers vary in size and generally are designed to support Maximum Transmission Unit (MTU)–sized frames.

After it determines the outbound interface for a packet, the forwarding processor does the following:

1. The forwarding processor sends a pointer with information about the packet (including its location in memory) to the virtual output queue of the outbound interface.
2. The line card's scheduler issues a request to the scheduler. The scheduler issues a grant, and the packet is sent from the buffer memory across the switching fabric to the outbound line card.
3. The outbound line card buffers the packets.
4. The L3 processor and associated Application–Specific Integrated Circuits (ASICs) on the outbound LC transmit the packet out the interface.

If the outbound interface is oversubscribed, it begins to buffer the excess packets. During periods of sustained oversubscription, the outbound LC's transmit queues fill. In this condition, the following happens depending on the outbound LC:

Engine Type of outbound LC	Response to Outbound	Error Counter
Engine 0 and 1	Sends congestion backpressure signal. The inbound interface begins to buffer the excess packets.	Ignored errors in the <b>show interfaces</b> command output and/or no mem drops in <b>execute-on slot &lt;slot#&gt; show controllers tofab QM stat</b> command output of the inbound LC, depending on its L3 Forwarding Engine. <sup>1</sup>
Engine 2, 3, 4	Drops any excess packets on the egress.	No mem drops in <b>execute-on slot &lt;slot#&gt; show controllers frfab QM stat</b> command output on the outbound LC.

<sup>1</sup>You will get ignored errors for L3 Engines 0, 1, and 2 ingress LCs. However, for four, 16 and more ports on Engine 2 LCs, the ignored counter will not increase.

On any intelligent networking device, when one or more high–speed interfaces feed a relatively low–speed interface, a mismatch in interface rates occurs. Since the slower–speed outbound interface cannot possibly

return buffers as fast as the faster inbound interface is sending them to the output hold queue, a delay in buffer return leads to some type of drops. This packet flow breaks the assumption that the outbound interface returns the buffer at the rate of buffer management time.

In addition to a mismatch in interface rates, ignored errors can increment when the rate of arriving packets is greater than the CPU can process them. This condition is very rare on the Cisco 12000 and usually results from a large number of very small packets, or when a CPU-intensive feature, such as Access Control Lists (ACLs) or traffic policing, is enabled on an LC that implements these features in software. This is the case for Engine 0 LCs where lots of features are implemented in software. However, on later engines, almost all the features are implemented in hardware. For instance, Engine 3 (IP Services Engine – ISE) and Engine 4+ line cards are designed for Edge applications and implement enhanced IP services (such as Quality of Service – QoS) in hardware with no performance impact. Examples of this hardware include 1–Port CHOC–48 ISE, 4–Port CHOC–12 ISE, 16–Port OC–3 POS ISE, 4–Port OC–12 POS ISE, 1–Port OC–48 POS ISE, and 1–Port OC–48 POS ISE.

The ignored counter can also be incremented whenever a packet arrives on an ingress line card and an appropriate size packet buffer is not available to handle this packet. However, this condition is very rare and it is not covered in this document.

## Managing Outbound Interface Oversubscription

The solution to ignored errors and no mem drops caused by output interface oversubscription is the same for any L3 Engine type — prevent buffer starvation. In other words, we need a mechanism that prevents the FrFab queues from filling.

### Engines 0 and 1

Simply put, the ignored counter is incremented when a packet arrives on an ingress line card (LC) and an appropriate size packet buffer is not available to handle this packet. Thus, ignored packets typically do not point to a bug in Cisco IOS software.

Here's a sample output from the **show interfaces** command with a non-null ignored counter on a Cisco 12000 Series Router:

```
router#show interfaces G3/0
GigabitEthernet3/0 is up, line protocol is up
  Hardware is GigMac GigabitEthernet, address is 0030.71f5.7980 (bia 0030.71f5.7980)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec, rely 255/255, load 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not set
  Full-duplex mode, link type is force-up, media type is SX
  output flow-control is unsupported, input flow-control is unsupported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:00, output hang never
  Last clearing of "show interface" counters 00:00:07
  Queueing strategy: fifo
  Output queue 0/40, 0 drops; input queue 0/75, 0 drops
  5 minute input rate 99000 bits/sec, 74 packets/sec
  5 minute output rate 104000 bits/sec, 68 packets/sec
    478 packets input, 71057 bytes, 0 no buffer
    Received 19 broadcasts, 0 runts, 0 giants, 0 throttles
    2 input errors, 2 CRC, 0 frame, 0 overrun, 25 ignored

    !--- Ignored counter is > 0. Ensure it is incrementing.

  0 watchdog, 53 multicast, 0 pause input
  541 packets output, 139133 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 pause output
0 output buffer failures, 0 output buffers swapped out
```

When the outbound LC is an Engine 0 or 1, it sends a backpressure message to the other LCs telling them to no longer send data to that particular LC. The inbound interface then buffers the excess packets in its ToFab queues corresponding to this destination slot.

To isolate the most likely cause of why the Ignored counter is increasing, you need to look at the ToFab queues of the ingress LC. You can either attach to the LC over the Maintenance BUS (MBUS) using the **attach** command, or use the **execute-on slot <slot#> show controllers tofab queue** command to check the ToFab queues. Execute this command a few times and look for the following symptoms:

- A decreasing and low value or value of 0 in the #Qelem column of a non-IPC free queue
- A large value in the #Qelem column in a destination slot queue.

## Engines 2, 3, 4

Line cards using a more recent L3 Engine architecture do not use a backpressure mechanism. Instead, when the interface is oversubscribed and a FrFab queue becomes depleted, the packets are simply dropped as they arrive on the egress line card.

Engine 2 LCs do not fall back to the next larger buffer pool when a smaller pool becomes depleted. The fall back mechanism has only been implemented for Engine 2 LCs on the ToFab side (Rx). If it happens, the "bump count" counter will increase in the output of the **execute-on slot <slot#> show controller tofab QM stat** command.

These drops are counted as *no mem drops* in the output of the **execute-on slot <slot#> show controllers frfab QM stat** command, as illustrated below:

```
Router#execute-on slot 1 show controller
frfab QM stat
===== Line Card (Slot 1) =====

174 no mem drop, 0 soft drop, 0 bump count

!--- Look for an incrementing value for the "no mem drop" counter

0 rawq drops, 0 global red drops, 0 global force drops
0 no memory (ns), 0 no memory hwm (Ns)
no free queue
0      0      0      0
0      0      0      0
0      0      0      0
0      0      0      0
0 multicast drops
Tx Counts
Interface 0
8390658710246 TX bytes, 2098330790 TX pkts, 212452 kbps, 6641 pps
Interface 1
0 TX bytes, 0 TX pkts, 0 kbps, 0 PPS Interface 2
0 TX bytes, 0 TX pkts, 0 kbps, 0 PPS Interface 3
0 TX bytes, 0 TX pkts, 0 kbps, 0 PPS
```

You need to find a way to prevent the FrFab side from buffering to the point where the LC either backs up to the inbound interface or simply drops the packets.

A simple solution for all the line cards, except Engine 2 LCs, is to reduce the number of buffers available to a particular outbound interface on a multi-interface LC. By default, an interface can use all of the carved FrFab buffers. Use the **tx-queue-limit** command to configure a non-default value. This prevents the egress LC from buffering more than the configured number of packets on the interface queue for that specific port. Make sure that you configure this number low enough so that it does not contain all the FrFab queues for this interface. Note that this method does not differentiate between high and low priority packets and simply implements tail drop more aggressively for a particular interface.

Engine 3 line cards require the use of Modular QoS CLI (MQC) instead of the legacy Command Line Interface (CLI). This command is not supported on Engine 2-based line cards.

Here's a configuration example using the legacy Class of Service (CoS) configuration:

```
interface POS 0/0
  tx-queue-limit <max Q length in packets>
```

Here's a configuration example using the MQC:

```
policy-map TX_QUEUE_LIMIT
  class class-default
    queue-limit

interface POS 0/0
  service-policy out TX_QUEUE_LIMIT
```

Another solution is to implement a faster output interface, which gives us a larger pipe. But larger pipes can fill quickly. Thus, the recommended solution is to implement Quality of Service (QoS) mechanisms on the outbound LC.

Cisco's Weighted Random Early Detection (WRED) feature implements a differentiated or intelligent drop mechanism. It is designed to work with adaptive traffic, such as TCP flows. It monitors the queue size and works to maintain a consistent average queue size by dropping packets randomly from various flows as the calculated average queue rises above a configurable minimum threshold.

When implemented on the Cisco 12000 Series, WRED can prevent the FrFab queues from filling and importantly is selective about which packets it drops. Engine 0 LCs support WRED in software, whereas Engine 1 LCs do not support WRED at all. The other L3 Engine LCs support WRED in hardware.

For more information on configuring WRED, see the following documents:

- Weighted Random Early Detection on the Cisco 12000 Series Router
- Configuring MPLS Classes of Service on a Cisco 12000 Series GSR Router

This congestion avoidance mechanism works only in a TCP-based environment. TCP responds appropriately – even robustly – to traffic drops by slowing down its traffic transmission. See [How TCP Handles Traffic Loss](#) and [How the Router Interacts with TCP](#) for details about how TCP reacts to packet loss.

Another supported QoS mechanism on the Cisco 12000 Series is traffic policing using Committed Access Rate (CAR) on Engine 0 and Engine 1 LCs, and a modified version of CAR known as Per Interface Rate Control (PIRC) on Engine 2 LCs. Configure traffic policing on the outbound interface.

## Managing Overloaded CPU on the Inbound Line Card

This situation is very rare!

You can check if the CPU is overloaded on the incoming LC using the **execute-on slot <slot#> show controllers tofab queues** command. If you see a very large number in the #Qelem column of the "Raw Queue" row, it means that too many packets are intended to be handled by the CPU (which is located on the LC itself). You will start to get ignored packets because the CPU cannot keep up with the amount of packets. These packets are directed to the CPU of the LC, not to the Gigabit Route Processor (GRP)!

What you need to do at this time is shift a part of the traffic from this inbound LC so that its CPU is less impacted.

You should also have a look at the LC configuration to check if there are some features configured on it that impact the CPU. Some features (such as CAR, ACL, and NetFlow) can degrade the performance of the LC when implemented in software (only on Engine 0 LCs). If this is the case, you should act accordingly by either removing the feature or upgrading the Cisco IOS software to a later release where the same feature implementation is improved (such as Turbo ACL). See Release Notes of the Cisco 12000 Series Routers to find out which features have been implemented or improved for different LCs.

Finally, the only solution may be to swap the LC for a more recent one where the requested feature is implemented in the hardware. This really depends on the Engine type of the LC.

You can use the following shortcut command to determine the L3 Engine type of an LC:

```
Router#show diag | i (SLOT | Engine)
...
SLOT 1 (RP/LC 1 ): 1 port ATM Over SONET OC12c/STM-4c Multi Mode
  L3 Engine: 0 - OC12 (622 Mbps)
SLOT 3 (RP/LC 3 ): 3 Port Gigabit Ethernet
  L3 Engine: 2 - Backbone OC48 (2.5 Gbps)
...
```

**Note:** Engine 3 (IP Services Engine – ISE) and Engine 4+ line cards are designed for Edge applications and implement enhanced IP services (such as QoS) in hardware with no performance impact.

## Summary of Design Guidelines

- Use Turbo ACLs, which optimize performance by allowing the router to compile the ACLs before downloading them to the LC processor.
- Avoid using the "log" keyword on ACLs.
- Avoid outgoing ACLs when possible. In a system with Engine 0, 1 and 2 LCs, all processing of ACLs is done on the inbound LC. Even outbound ACL filtering is done on the inbound card once it knows to which outbound interface the packet is destined. For this reason, configuring an outbound ACL on an interface affects all LCs in the system. In addition, Engine 2 LCs can perform either incoming or outgoing ACLs, but not both simultaneously in the ASIC that performs hardware-forwarding. If you configure both inbound and outbound ACLs, the LC falls back to CPU-based forwarding for outgoing access lists, impacting the switching performance of the LC. However, newer Engines such as Engine 3 and Engine 4+ are highly optimized for enhanced IP services like ACLs and process outbound ACLs on the outgoing LC.
- Assign traffic requiring specific features to one set of LCs.
- Assign traffic not requiring features to another set of LCs to maintain peak packet forwarding performance.
- Use LCs with higher Engine types when high performance is needed.
- Design backbone- or core-facing LCs to run features supported in hardware or microcode.

# Case Study

This case study shows how to troubleshoot incrementing ignored errors on an interface of an LC in slot 6.

## Step 1 – Check the ToFab queues in slot 6.

```
Router#exec slot 6 show controllers tofab queue
===== Line Card (Slot 6) =====
Carve information for ToFab buffers
  SDRAM size: 134217728 bytes, address: 30000000, carve base: 30019100
  134115072 bytes carve size, 4 SDRAM bank(s), 8192 bytes SDRAM pagesize, 2 carve(s)
  max buffer data size 4544 bytes, min buffer data size 80 bytes
  174538/174538 buffers specified/carved
  110797216/110797216 bytes sum buffer sizes specified/carved
    Qnum      Head      Tail          #Qelem  LenThresh
    ----      -
4 non-IPC free queues:
  88964/88964 (buffers specified/carved), 50.97%, 80 byte data size
  1          21120          84604      81074     262143
  54076/54076 (buffers specified/carved), 30.98%, 608 byte data size
  2          122270         116965     49567     262143
  26165/26165 (buffers specified/carved), 14.99%, 1568 byte data size
  3          164160         145355     19518     262143

  !-- Out of the 26165 buffers that are carved, only 19518 are available

  5233/5233 (buffers specified/carved), 2.99%, 4544 byte data size
  4          172325         172088     5233      262143
IPC Queue:
  100/100 (buffers specified/carved), 0.5%, 4112 byte data size
  30         61           60           100       262143
Raw Queue:
  31         44229          88895        0         43634

  !-- The Raw Queue has a low or 0 value for the #Qelem column, indicating
  !-- that the CPU is not overwhelmed with packets destined to it.

ToFab Queues:
  Dest
  Slot
  0          73769          60489        0         262143
  1          7909           27395        0         262143
  2          61416          71346        0         262143
  3          80352          14567        0         262143
  4          138236         107121       18955     262143

  !-- 18955 packets are waiting for space in the outbound queues
  !-- on the LC in slot 4.

  5          4852           48171        0         262143
  6          98318          111757       0         262143
  7          44229          88895        0         262143
  8          0              0            0         262143
  9          0              0            0         262143
  10         0              0            0         262143
  11         0              0            0         262143
  12         0              0            0         262143
  13         0              0            0         262143
  14         0              0            0         262143
  15         0              0            0         262143
Multicast   0              0            0         262143
```

## Step 2 – Check the FrFab queues in slot 4.

Since the ToFab queue output indicated a large number of queued packets destined for the LC in slot 4, check the FrFab queues on this LC.

```
Router#exec slot 4 show controllers frfab queue
===== Line Card (Slot 4) =====
Carve information for FrFab buffers
  SDRAM size: 67108864 bytes, address: 20000000, carve base: 2002D100
  66924288 bytes carve size, 0 SDRAM bank(s), 0 bytes SDRAM pagesize, 2 carve(s)
  max buffer data size 4544 bytes, min buffer data size 80 bytes
  65534/65534 buffers specified/carved
  66789056/66789056 bytes sum buffer sizes specified/carved
    Qnum      Head      Tail          #Qelem  LenThresh
    ----      -
4 non-IPC free queues:

  26174/26174 (buffers specified/carved), 39.93%, 80 byte data size
  1          10123   4332          14515   65535

  19630/19630 (buffers specified/carved), 29.95%, 608 byte data size
  2          27898   37167         12279   65535

  13087/13087 (buffers specified/carved), 19.96%, 1568 byte data size
  3  0          52275         0        65535

  !-- Zero buffers available for this pool

  6543/6543 (buffers specified/carved), 9.98%, 4544 byte data size
  4          60805   60804         6543    65535

IPC Queue:

  100/100 (buffers specified/carved), 0.15%, 4112 byte data size
  30       75       74           100     65535

Raw Queue:

  31        0        80           0       65535

Interface Queues:

  0         0        39413         0       65535
  1         0        44192         0       65535
  2        48426    58230        32111    65535

  !-- Interface 2 is using half or 32111 of the carved packet buffers

  3         0        41219         0       65535
```

## Step 3 – Compare the output of the show interfaces command for the same oversubscribed interface.

Match the oversubscribed interface indicated in the **show controllers frfab queue** output with the **show interfaces** output for the same interface. The following output confirms that the output interface rate is at the line rate and is oversubscribed:

```
Router#show interfaces POS 4/2
POS4/2 is up, line protocol is up
  Hardware is Packet over SONET
  Description: Pacbell OC3 to other ISP...
```

```
Internet address is 10.10.10.10/30
MTU 4470 bytes, BW 155000 Kbit, DLY 100 usec, rely 255/255, load 156/255
Encapsulation HDLC, crc 32, loopback not set
Keepalive set (10 sec)
Scramble enabled
Last input 00:00:01, output 00:00:03, output hang never
Last clearing of "show interface" counters never
Queueing strategy: FIFO  Output queue 0/300, 0 drops; input queue 0/300, 0 drops
5 minute input rate 20274000 bits/sec, 6263 packets/sec
5 minute output rate 148605000 bits/sec, 28776 packets/sec

!-- The output interface rate is at line rate which means that the interface is
!-- oversubscribed.
```

```
1018621328 packets input, 2339977099 bytes, 0 no buffer
Received 0 broadcasts, 1 runts, 0 giants, 0 throttles
      0 parity
1 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
378645 packets output, 156727974 bytes, 0 underruns
0 output errors, 0 applique, 0 interface resets
0 output buffer failures, 0 output buffers swapped out
1 carrier transitions
```

## Step 4 – Implement a solution.

See the solutions sections of this document for the next steps to resolve incrementing ignored errors based on the architecture of the particular outbound interface. For example, on an Engine 0 LC, try diverting some traffic to another interface or, as a temporary measure, reduce the number of packet buffers that this particular interface can use from the line card's free queues. Use the following command:

```
Router(config)#int POS 4/2
Router(config-if)#tx-queue-limit 5000
```

## Cisco IOS Software Bugs

Sometimes counters increment due to a Cisco IOS software defect. Be sure you are running the latest available Cisco IOS software release in your train to get rid of all the bugs that have already been fixed. If you still see ignored packets, and the information in this document does not solve your issue, contact Cisco's Technical Assistance Center (TAC) for assistance.

---

## Related Information

- [Troubleshooting Input Drops on the Cisco 12000 Series Internet Router](#)
  - [How To Read the Output of the show controllers frfab | tofab queue Commands](#)
  - [Weighted Random Early Detection on the Cisco 12000 Series Internet Router](#)
  - [Configuring MPLS Classes of Service on a Cisco 12000 Series GSR Router](#)
  - [How TCP Handles Traffic Loss](#)
  - [How the Router Interacts with TCP](#)
  - [Configuring Committed Access Rate](#)
  - [Release Notes of the Cisco 12000 Series Internet Routers](#)
  - [Technical Support – Cisco Systems](#)
-

