

Troubleshooting Input Queue Drops and Output Queue Drops

Document ID: 6343

Interactive: This document offers customized analysis of your Cisco device.

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

Processing and Switching

Input Queue Drops

- Troubleshoot Input Queue Drops

Output Queue Drops

- Troubleshoot Output Queue Drops

Commands to Obtain More Information

- show interfaces switching

- show interfaces stats

- ip accounting mac-address

- show interfaces mac-accounting

Related Information

Introduction

This document discusses input and output queue drops taken from the output of the **show interfaces** command on the router. This document describes what these drops mean, the type of problems they indicate, and how to troubleshoot the source of these problems. It provides some tips on how to prevent these problems.

Note: Drops can often be useful, because they trigger the flow control mechanisms of upper layer protocols (for example, drops decrease the TCP window size).

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

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.

Processing and Switching

In IP networks, routers make forwarding decisions based on the contents of the routing table. When a router searches the routing table, it looks for the longest match for the destination IP address. The router does this at the process level. Therefore, the search process is queued among other CPU processes, because of which, the lookup time is unpredictable and can be very long. Therefore, a number of switching methods based on exact-match-lookup have been introduced in Cisco IOS® Software.

The main benefit of exact-match-lookup is that the lookup time is deterministic and very short. This has significantly shortened the time a router takes to make a forwarding decision. Therefore, routines that perform the search can be implemented at the interrupt level. This means, the arrival of a packet triggers an interrupt, which causes the CPU to postpone other tasks and handle the packet. The legacy method to forward packets is to look for a best match in the routing table. This cannot be implemented at interrupt level and must be performed at process level. For a number of reasons, some of which are mentioned in this document, the longest-match-lookup method cannot be completely abandoned, so these two lookup methods exist in parallel on Cisco routers. This strategy has been generalized, and is now also applied to IPX and AppleTalk.

For more information on Cisco IOS Software switching paths, refer to Performance Tuning Basics.

Input Queue Drops

When a packet enters the router, the router attempts to forward it at interrupt level. If a match cannot be found in an appropriate cache table, the packet is queued in the input queue of the incoming interface to be processed. Some packets are always processed, but with the appropriate configuration and in stable networks, the rate of processed packets must never congest the input queue. If the input queue is full, the packet is dropped.

Here is a sample output:

```
router#show interfaces ethernet 0/0
...
Input queue: 30/75/187/0 (size/max/drops/flushes); Total output drops: 0
Output queue :0/40 (size/max)...
```

In this sample output, there is no way to see exactly which packets have been dropped. In order to troubleshoot input queue drops, you must find out which packets fill the input queue. In this example, 30 packets are in the input queue of interface ethernet0/0 when the **show interfaces ethernet 0/0** command is issued. The queue depth is 75 packets and there have been 187 drops since the interface counters were last cleared.

The system counts input queue drops if the number of packet buffers allocated to the interface is exhausted or reaches its maximum threshold. You can increase the maximum queue value with the **hold-queue <value>** command for each interface (the queue length value can be between 0 and 4096. The default value is 75).

Note: Shared-memory routers (1600, 2500, and 4000 series), also use the input queue for fast-switched traffic. If you get input queue drops on those platforms, ensure that all traffic uses the best switching path available (see Performance Tuning Basics). Input queue drops generally occur when a packet is process-switched. A process switch means that the router cannot use a preferable route-cache method, such as fast switching or Cisco Express Forwarding (CEF), to handle the forwarding decision. If input drops are

still present, it implies that there is simply too much traffic. Consider a hardware upgrade, or try to decrease the traffic load.

These are the conditions for input queue drop counter. They usually occur when the router receives bursty traffic and cannot handle all packets.

- The Rx FIFO which is accessible by the interface PHY and interface DMA is full and any new frames that arrive in this condition will be dropped (normally called as overflow) and the rx_overflow counter (seen through **show controller interface-id**) will be incremented. When rx_overflow counter is incremented by one, it indicates that overflow condition has occurred once and is not indicative of the number of frames dropped.
- The Rx ring which is accessible by the interface DMA and interface driver code is full. Any new frame transfers from the DMA cannot proceed with this condition, since there are no free entries in Rx ring and hence the frames sent are dropped (termed as overrun condition). The rx_int_drop counter (seen through **show controller interface-id**) is also incremented by one. Again, if rx_int_drop is incremented by one it indicates that there is one occurrence of an overrun condition, and the number of frames dropped is not known.

The input hold queue size can be increased from the default 75 packets. The hold queue stores packets received from the network that wait to be sent to the client. Cisco recommends that the queue size not exceed ten packets on asynchronous interfaces. For most other interfaces, queue length must not exceed 100. The input hold queue prevents a single interface from flooding the network server with too many input packets. Further input packets are discarded if the interface has too many input packets outstanding in the system.

```
Router(conf-if)# hold-queue length in
```

For Catalyst Switches, Cisco recommends to make this adjustment on all L3 interfaces on the device, both physical interfaces and VLAN interfaces. L2 ports configured with the **switchport** command can be left at the default value.



Caution: An increase in the hold queue can have detrimental effects on network routing and response times. For protocols that use SEQ/ACK packets to determine round-trip times, do not increase the output queue. Dropping packets instead informs hosts to slow down transmissions to match available bandwidth. This is generally better than duplicate copies of the same packet within the network, which can happen with large hold queues.

Troubleshoot Input Queue Drops

You can successfully troubleshoot input queue drops while packets constantly arrive in the input queue. You cannot troubleshoot a congestion that occurred in the past. If more than one routed protocol is configured on the interface, first determine the protocol that congests the input queue. Here is the fastest way to do this is:

1. Determine the suspect protocol. Check the CPU utilization in `<protocol> Input` processes. To do so, run the **show processes cpu** exec command. If Cisco IOS Software version 12.1 or higher currently runs on the router, you can shorten the output of the **show processes CPU** command through the output modifiers:

```
router#show processes CPU | i ^PID|Input
PID  Runtime(ms)  Invoked  uSecs   5Sec   1Min   5Min  TTY  Process
 10      8503       1713    4963   0.00%  0.00%  0.00%  0  ARP Input
 24     69864      11429   6112   0.08%  0.11%  0.10%  0  Net Input
 28    55099       8942   6161  26.20% 20.07% 19.26%  0  IP Input
 37         4         2     2000   0.00%  0.00%  0.00%  0  SSCOP Input
 40         8         2     4000   0.00%  0.00%  0.00%  0  ILM Input
 49         8         1     8000   0.00%  0.00%  0.00%  0  Probe Input
```

| | | | | | | | |
|-----|-------|------|------|-------|-------|-------|----------------|
| 50 | 28209 | 4637 | 6083 | 0.00% | 0.03% | 0.04% | 0 RARP Input |
| 59 | 8 | 2 | 4000 | 0.00% | 0.00% | 0.00% | 0 SPX Input |
| 61 | 8 | 2 | 4000 | 0.00% | 0.00% | 0.00% | 0 Tag Input |
| 68 | 20803 | 3392 | 6132 | 0.00% | 0.03% | 0.00% | 0 IPX Input |
| 104 | 4 | 1 | 4000 | 0.00% | 0.00% | 0.00% | 0 IPXWAN Input |
| 107 | 8 | 1 | 8000 | 0.00% | 0.00% | 0.00% | 0 AT Input |

Table 1 lists the possible input processes and types of packets that can congest the input queue:

| Input Process That Uses CPU Cycles | Type of Packets |
|------------------------------------|-----------------|
| IP | IP |
| AT | AppleTalk |
| IPX, SPX, or IPXWAN | IPX |
| ARP | IP ARP |

Other input processes are not likely to congest the input queue.

- Find out whether packets that congest the input queue are destined for the router, or are forwarded through the router. Run the **show interfaces [type number] switching** command from exec mode.

Note: The **show interfaces [type number] switching** command is hidden, and does not show up if you use the "?" or TAB keys on the command line interface. Type the full command on the router. This command is not documented in the Command Reference Guide

```
router#show interfaces ethernet 0/0 switching
Ethernet0/0
...
Protocol          Path      Pkts In   Chars In   Pkts Out   Chars Out
...
IP                Process   12142    2211929    35         5169
                Cache misses 10212
```

Check whether the number of processed packets received is followed by a high number of cache misses. If so, this indicates that the packets, which congest the input queue, are forwarded through the router. Otherwise, these packets are destined for the router.

- If packets are destined for the router, find out which higher-layer protocol congests the input queue. For this, use one of these **show traffic** exec commands:

- ◆ **show ip traffic**
- ◆ **show ipx traffic**
- ◆ **show appletalk traffic**

Note: These commands are applicable only if you suspect any of the input processes listed in table 1.

- Try to get more information about the packets that congest the input queue. For this, you must debug the received packets. The previous steps indicate the debug commands that you need to enable.

Note: You can execute this directly, even if you do not perform the previous steps. However, when you debug, several messages are generated, and they can be hard to read. When you follow all the previous steps, you get an indication of what to look for in the debug output.



Warning: Debug with extreme caution. Otherwise, CPU utilization can increase considerably.

Do not turn debugging on for more than 5 to 10 seconds. For more information on how to use the debug commands, refer to Using Debug Commands. Never disable console logs, terminal logs, and logs on a syslog server. Enable buffer logs, and increase the logging buffer size. A good value for the

size of logging buffer would be 128000 bytes. Use these commands:

- ◆ **no logging <host>**
- ◆ **logging buffered 128000 debugging**

The output must be sufficient to locate the source of the problem. You can check the debug output with the **show log** command after you complete the debug session. Table 2 lists the **debug** commands to issue based on the type of packets that congest the input queue:

| Type of Packets that Congest the Input Queue | Debug Command to be Used |
|--|-------------------------------|
| IP | debug ip packet |
| AppleTalk | debug apple packet |
| IPX | debug ipx packet |
| ARP | debug arp |

For more information, refer to the Cisco IOS Debug Command Reference.

Alternatively, you can use the **show buffers input-interface** *[interface type]* *[interface number]* **header** command to find out the types of that packets fill up the input queue.

Note: This is only useful if there are a lot of packets in the input queue.

```
Router#show buffers input-interface serial 0/0
Buffer information for Small buffer at 0x612EAF3C
  data_area 0x7896E84, refcount 1, next 0x0, flags 0x0
  linktype 7 (IP), enctype 0 (None), encsize 46, rxtyp 0
  if_input 0x6159D340 (FastEthernet3/2), if_output 0x0 (None)
  inputtime 0x0, outputtime 0x0, oqnumber 65535
  datagramstart 0x7896ED8, datagramsize 728, maximum size 65436
  mac_start 0x7896ED8, addr_start 0x7896ED8, info_start 0x0
  network_start 0x7896ED8, transport_start 0x0
  source: 212.176.72.138, destination: 212.111.64.174, id: 0xAAB8,
  ttl: 118, prot: 1
Buffer information for Small buffer at 0x612EB1D8
  data_area 0x78A6E64, refcount 1, next 0x0, flags 0x0
  linktype 7 (IP), enctype 0 (None), encsize 46, rxtyp 0
  if_input 0x6159D340 (FastEthernet3/2), if_output 0x0 (None)
  inputtime 0x0, outputtime 0x0, oqnumber 65535
  datagramstart 0x78A6EB8, datagramsize 728, maximum size 65436
  mac_start 0x78A6EB8, addr_start 0x78A6EB8, info_start 0x0
  network_start 0x78A6EB8, transport_start 0x0
  source: 212.176.72.138, destination: 212.111.64.174, id: 0xA5B8,
  ttl: 118, prot: 1
```

Most of the time, one type of packet is present in large quantities. Here, for example, there are several Internet Control Message Protocol (ICMP) packets (IP protocol 1).

If the problem is an incorrect router configuration (for example, both fast switching and Cisco express forwarding (CEF) are disabled), there is probably no particular pattern in the debugs, or in the output of the **show buffers input-interface** command.

5. When you have determined the type of packets that congest the input queue, the next step is to check whether you can prevent this congestion.

There are several reasons why packets must be processed:

- ◆ **Improper router configuration** Switching paths that operate at the interrupt level are disabled on relevant interfaces.

To check which switching paths are configured on an interface, run the **show <protocol> interface [type number]** command.

- ◇ In order to enable legacy fast switching, configure it on output interfaces.
- ◇ In order to enable netflow switching, configure it on input interfaces.
- ◇ In order to enable Cisco express forwarding (CEF), you have to enable CEF globally (on the entire router) and locally (on the incoming interface).

For more information, see Cisco IOS Switching Services Configuration Guide.

- ◆ **Local destination** Packets are destined for the router.

- ◇ In stable networks, the number of routing updates must not be excessive. In unstable networks, frequent updates of large routing tables can congest the input queue.
- ◇ Check whether excessive traffic is directed to the router itself (with, for example, Simple Network Management Protocol (SNMP), telnet, Trivial File Transfer Protocol (TFTP), and ping). Debug the packets for the relevant protocol to identify the source of these packets. When you find the source, eliminate it.

- ◆ **Reliable Open System Interconnection (OSI) layer 2 protocol is used for transport** Packets that go through serial interfaces with X.25 encapsulation must be processed because in the X.25 protocol suite, flow control is implemented on the second OSI layer.
- ◆ **Software compression** If the packet comes in or has to be forwarded through an interface on which software compression is configured, the packet has to be processed.
- ◆ **Other features are unsupported at interrupt level** This is highly dependent on the Cisco IOS Software release that runs on the router. Check the release notes to see which features are supported at interrupt level. For example, in earlier Cisco IOS Software versions, multilink PPP packets had to be processed. In higher Cisco IOS Software versions, they can be fast-switched or even CEF-switched. Features such as encryption, local-area transport (LAT) translation, and data-link switching plus (DLSW+) are not yet fast-switched.
- ◆ **Excessive traffic through the router, where each packet header intentionally contains different information** Based on the configured switching path, the first packets to a destination, or in a flow, are always processed. This is because, there are no entries in the cache that matches them. If a device sends packets at an extremely high rate, and there is no match in the cache, those packets can congest the input queue.

The source of these packets are revealed after the debug session. If the source address is always different, you must continue to troubleshoot on the upstream device, from which the packet is received. If the interface on the router is connected to a broadcast medium, you can determine the Media Access Control (MAC) address of the source or the upstream device:

Configure MAC accounting on the interface with the **ip accounting mac-address input** interface configuration command. After that, issue the **show interfaces mac-accounting exec** command. This command reveals the MAC address that has sent the packets at an excessive rate.

Output Queue Drops

Output drops are caused by a congested interface. For example, the traffic rate on the outgoing interface cannot accept all packets that should be sent out. The ultimate solution to resolve the problem is to increase the line speed. However, there are ways to prevent, decrease, or control output drops when you do not want to increase the line speed. You can prevent output drops only if output drops are a consequence of short bursts of data. If output drops are caused by a constant high-rate flow, you cannot prevent the drops. However, you can

control them.

When packets are processed, they are sent to the output queue of the outgoing interface. Issue the **show interfaces** exec command to view the size of the queue, the current number of packets in the queue, and the number of drops. Based on the type of interface and the type of queueing configured, the number of output queue drops is not explicitly shown, because the output drops counter summarizes the output drops separately at the processing level and at the interrupt level:

```
router#show interfaces serial 0/0
...
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: weighted fair
Output queue: 0/1000/64/0 (size/max total/threshold/drops)
...
router#show interfaces serial 0/0
...
Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
Queueing strategy: fifo
Output queue :0/40 (size/max)
...
```

However, it takes longer to process a packet than to send the packet from the output queue to the wire. Therefore, it is highly unlikely that output queue drops (drops at processing level) can occur without drops at interrupt level. Output queue drops occur only if the interface is already congested at interrupt level, so that packets cannot be pulled out of the output queue before the queue becomes full. Therefore, output drops at processing level (output queue drops) and output drops at interrupt level always occur together, and there is practically no need to distinguish between these two counters.

Note: However, there is one exception. If the output queue is constantly full and if no packets are sent out of the interface at all, you must check for a hardware failure on the interface.

Troubleshoot Output Queue Drops

You can decrease, or even prevent, output drops if you adjust the configuration of these features:

- **Duplex mode** If the interface works in half-duplex mode, configure it (if possible) to work in full-duplex.
- **Layer 2 windowing mechanism** If x.25 encapsulation is configured on the interface, increase the x.25 window size. For more information, see Setting Default Window Sizes.
- **Distributed switching** On Cisco 7500 routers, if Versatile Interface Protocol (VIP) cards are installed in the chassis, enable distributed switching. When you do so, the incoming VIP buffers up to 1 second of traffic for the interface if the outgoing interface is congested. This is called **rx-side buffering**.

Note: Never increase the output queue in an attempt to prevent output drops. If packets stay too long in the output queue, TCP timers can expire and trigger retransmission. Retransmitted packets only congest the outgoing interface even more.

If output drops still occur after you adjust the configuration of the router as recommended, it means that you cannot prevent or decrease output drops. However, you can control them, and this can be as effective as prevention. There are two approaches to control output drops:

- Congestion management
- Congestion avoidance

Both approaches are based on traffic classification, and you can use them in parallel.

Congestion management ensures, with appropriate configuration, that important packets are always forwarded, while less important packets are dropped when the link is congested. Congestion management comprises fancy queueing mechanisms such as:

- Priority queueing
- Class-based weighted fair queueing

Congestion avoidance is based on intentional packet drops. The window size in TCP connections depends on the round trip time. Therefore, these intentional drops slow down the rate at which the source device sends packets. Congestion avoidance uses weighted random early detection.

If unwanted output drops still occur after you implement these mechanisms, you need to increase the line speed.

Commands to Obtain More Information

Here are some commands that provide more information about queue drops:

- **show interfaces switching**
- **show interfaces stats**
- **ip accounting mac-address**
- **show interfaces mac-accounting**

If you have the output of a **show interfaces** command from your Cisco device, you can use to display potential issues and fixes. To use, you must be a registered customer, be logged in, and have JavaScript enabled.

show interfaces switching

Description

This command shows the number of packets sent and received on an interface, classified based on the switching path. This is a hidden command.

Format

```
show interfaces [type number] switching
```

Sample Output

```
Ethernet0/0
      Throttle count          0
      Drops      RP          0      SP  0
      SPD Flushes      Fast  0      SSE  0
      SPD Aggress      Fast          0
      SPD Priority      Inputs  86      Drops  0
      Protocol      Path      Pkts In      Chars In      Pkts Out      Chars Out
      Other      Process      75          6728          79          4740
      Cache misses          0
      Fast          0          0          0          0
      Auton/SSE          0          0          0          0
      IP      Process      142          11929          35          5169
      Cache misses          0
      Fast          0          0          0          0
      Auton/SSE          0          0          0          0
      AppleTalk      Process      0          0          25          1635
      Cache misses          0
```

```

Fast 0 0 0 0
Auton/SSE 0 0 0 0
DEC MOP Process 0 0 2 154
Cache misses 0
Fast 0 0 0 0
Auton/SSE 0 0 0 0
ARP Process 56 3580 13 780
Cache misses 0
Fast 0 0 0 0
Auton/SSE 0 0 0 0
CDP Process 90 26906 27 8900
Cache misses 0
Fast 0 0 0 0
Auton/SSE 0 0 0 0

```

| Field | Definition |
|------------------------------|---|
| <protocol> Process | Number of processed packets. This includes packets destined for the router, and packets for which there is no entry in the appropriate switching cache table. |
| Cache misses | Packets that are forwarded through process level (for which there is no entry in fast switching cache). |
| Fast | Packets forwarded at interrupt level. |

show interfaces stats

Description

This command is similar to the **show interfaces switching** command, and provides information on the number of packets that are process-switched, fast-switched (any fast switching path), and distributed-switched (for VIP capable platforms). This is a hidden command.

Format

```
show interfaces [type number] stats
```

Sample Output

```

Router#show interfaces stats
FastEthernet8/0/0
Switching path  Pkts In  Chars In  Pkts Out  Chars Out
Processor        64      38646     323       32790
Route cache     477985  611343050 14815     18948150
Distributed cache 0         0         3564     4558356
Total           478049  611381696 18702     23539296

Serial12/0/0
Switching path  Pkts In  Chars In  Pkts Out  Chars Out
Processor        37       3783      36        2299
Route cache     14815    18800000 45118     59862772
Distributed cache 3450     4378520  0         0
Total           18302    23182303 45154     59865071

Interface Serial12/0/1 is disabled
...

```

ip accounting mac-address

Description

This is an interface configuration command. It accounts for the received or transmitted packets, classified based on the source or destination MAC address.

Format

`ip accounting mac-address {input/output}`

show interfaces mac-accounting

Description

This is an exec command. It shows the number of packets sent and received classified based on the destination and source MAC address.

Format

`show interfaces [type number] mac-accounting`

Sample Output

```
router#show interfaces ethernet 0/0 mac-accounting
Ethernet0/0
  Input(494 free)
    0000.0c5d.92f9(58 ): 1 packets, 106 bytes, last: 4038ms ago
    0004.c059.c060(61 ): 0 packets, 0 bytes, last: 2493135ms ago
    00b0.64bc.4860(64 ): 1 packets, 106 bytes, last: 20165ms ago
    0090.f2c9.cc00(103): 12 packets, 720 bytes, last: 3117ms ago
                        Total: 14 packets, 932 bytes
  Output (511 free)
    0090.f2c9.cc00(103): 8 packets, 504 bytes, last: 4311ms ago
                        Total: 8 packets, 504 bytes
```

Related Information

- [Performance Tuning Basics](#)
- [Input Queue Overflow on an Interface](#)
- [Output Queue Overflow on an Interface](#)
- [Troubleshooting Input Drops on the Cisco 12000 Series Internet Router](#)
- [Troubleshooting Output Drops with Priority Queueing](#)
- [Technical Support – Cisco Systems](#)

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

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jul 29, 2008

Document ID: 6343
