Troubleshoot Interface Packet Drops in IOS XE Routers

Contents

Introduction

This document describes how to troubleshoot interface packet drops on Cisco IOS® XE routers.

Requirements

Basic knowledge on packet flow and Cisco IOS XE.

Components Used

This document is based on Cisco IOS XE router platforms such as ISR 4000 and ASR1000 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, ensure that you understand the potential impact of any command.

Background Information

On Cisco IOS XE routers, packet drops can occur across different components, including:

- Shared Port Adapter (SPA) interfaces.

- Quantum Flow Processor (QFP): Dataplane processor that holds the PPEs (packet processor engines).

These drops can be observed in either the input or output direction on the interfaces. When analyzing the QFP, the primary focus is on output drops.

For packet drops affecting the control plane on Cisco IOS XE devices, refer to punt drops as outlined in the guide: Troubleshoot Policer Surpassed a Threshold.

Input Drops

Interfaces can experience input drops in the input queues. This counter can be seen with the command **show interfaces** in the input queue field, drops counter section:

```
---- show interfaces  ----

GigabitEthernet0/0/0 is up, line protocol is up
<snip>
  Input queue: 0/375/71966/0 (size/max/drops/flushes); Total output drops: 47009277
<snip>
```

The input drop counter on interfaces is also visible in the **show interface summary** command, under the **IQD** column, which represents packets dropped from the input queue.

```
---- show interface summary  ----

  Interface                   IHQ       IQD       OHQ       OQD      RXBS      RXPS      TXBS      TXPS
------------------------------------------------------------------------------------------------------
* Te0/0/0                     0         0         0         0  29544000      2830   1957000      1446
* Te0/0/1                     0         0         0         0  23476000      2555  16655000      3346
* GigabitEthernet0/0/0        0     71966         0  47019440  18852000      5321  59947000      6064
```

Input drops on interfaces typically occur when the input queues become overwhelmed and cannot be processed in time. As a result, packets can be selectively discarded based on the queuing algorithm in use.

Possible causes for having input drops include:

- The sending device is transmitting packets at a rate higher than what the receiving device can process, leading to input queue exhaustion
- Burst or microbursts in traffic patterns
- Platform limitations

You can try to increase the input queue size using command **hold-queue** under interface level:

```
Router(config-if)#hold-queue ?
  <0-240000>  Queue length
```

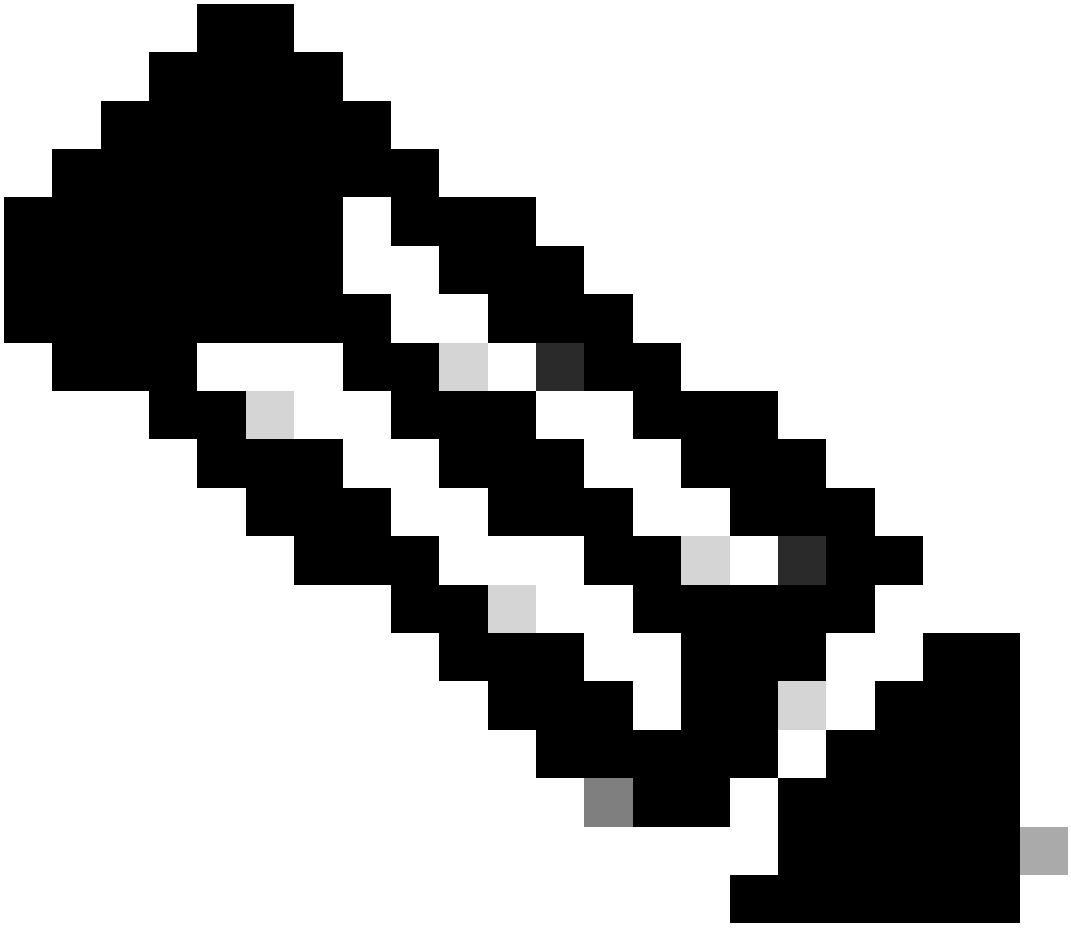

Note: **Hold-queue** command can not be effective on some platforms. Verify platform specifications or file a case with TAC.

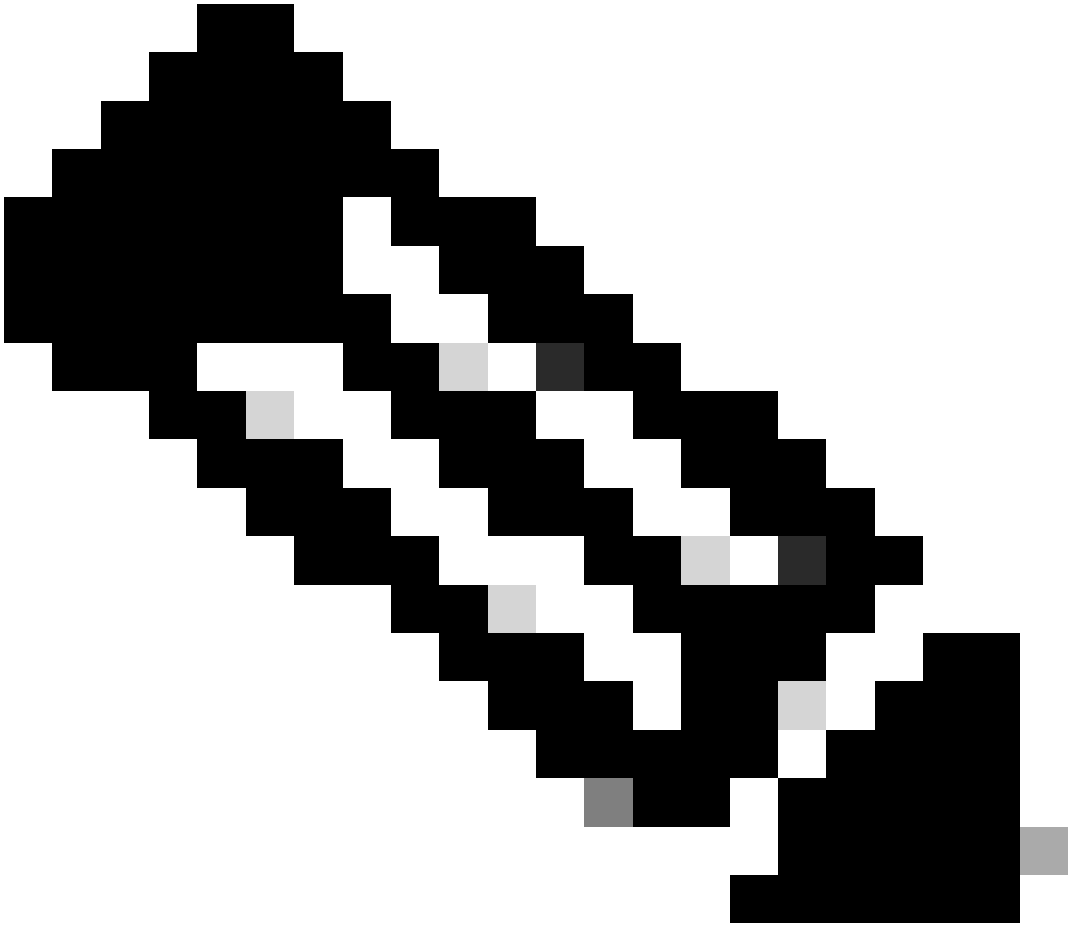

Note: Flow control mechanisms can also be used to send pause frames from the receiving device towards the sending device. Review more information about flow control within the Interface and Hardware Components Configuration Guide for the specific platform.

Output Drops on Interfaces

Output drops on interfaces manifest in the output queues and can be seen with the command **show interfaces**:

```
---- show int gi 1/0/46  ----

GigabitEthernet1/0/46 is up, line protocol is up (connected)
<snip>
  Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 154786
```

The Total Output Drops counter indicates saturation in the output queues of the affected interface. This

condition can be exacerbated by mechanisms such as Quality of Service (QoS), which can selectively discard packets to manage congestion.

Since QoS alters the priority of the traffic, another troubleshooting step is to verify whether the interface is using a non-default queueing strategy via a policy-map configured in the output direction using the command service-policy output.

```
interface GigabitEthernet0/1
  service-policy output PRIORITIZE-VOICE
```

To verify if the output drops are due to the quality of service mechanism implemented, use the command **show policy-map interface <interface-name> out**. This is shown in this example:

```
---- show policy-map interface gi0/0/0 output  ----

 GigabitEthernet0/0/0

  Service-policy output: PRIORITIZE-VOICE

    queue stats for all priority classes:
      Queueing
      queue limit 512 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0

    Class-map: VOICE (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match:  dscp ef (46)
      Priority: Strict, b/w exceed drops: 0


    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0000 bps, drop rate 0000 bps
      Match: any

      queue limit 4166 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 0/0
Router#
```

This command shows the drops due to quality of service mechanism among the classes configured.

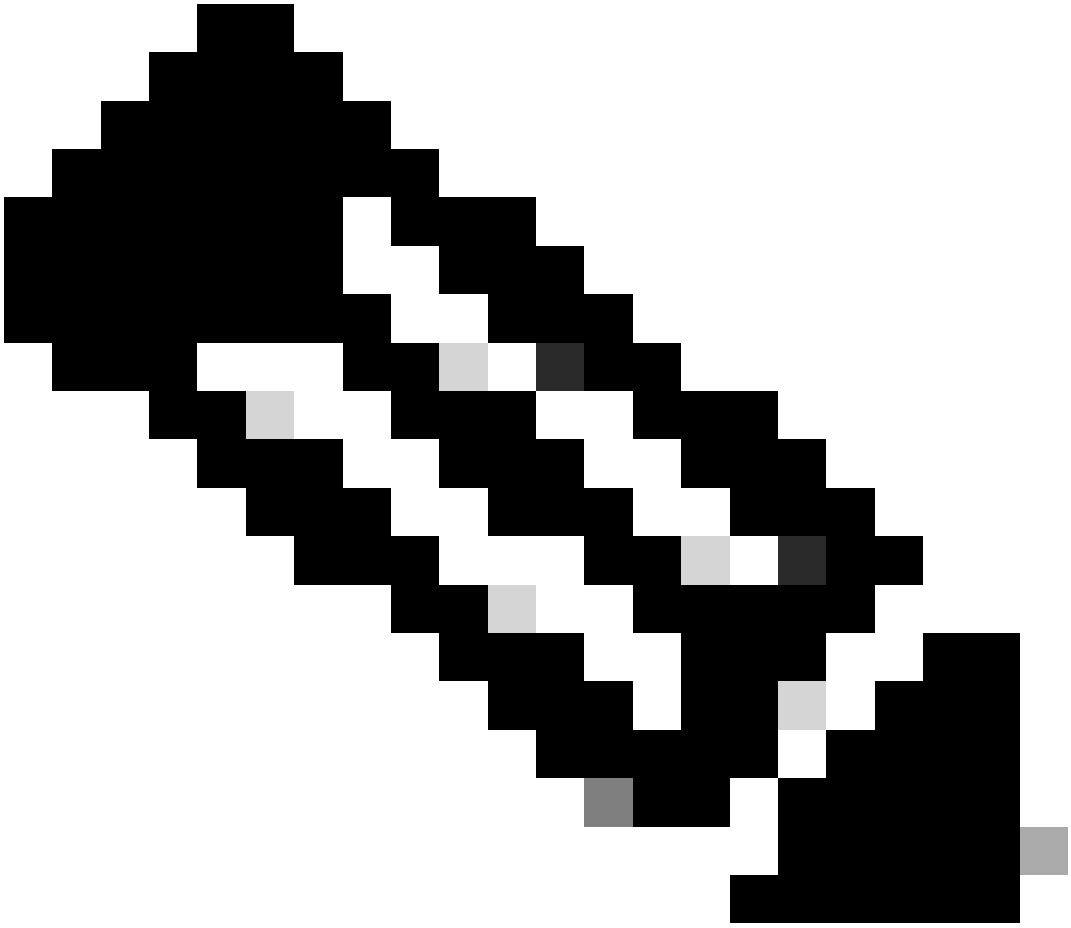

Note: If the output drops on the interface correlate with the drops seen in the policy map, the drop is generally expected due to the quality of service configured. Engage TAC if needed to dig deeper into the quality of service mechanism used and refer to the corresponding guides for this feature.

For additional information on QoS working and how can be implemented, please refer to the Quality of Service Configuration Guide, Cisco IOS XE 17.x configuration guide.

To view the queueing strategy, use the **show interfaces** command and check the queueing strategy value. By default, the outbound packet processing strategy is first-in, first-out (FIFO).

```
---- show interfaces gigabitEthernet 0/0/0  ----

<snip>
  Queueing strategy: Class-based queueing
```

Output Drops without Quality of Service in Place

If the interface does not have a policy-map associated in the output direction for quality of service, other causes can cause the output drops.

Some of the reasons for output drops on an interface that does not have quality of service are:

- Incoming interfaces forming port channels and oversubscribing single output interface
- Quantum flow processor (QFP) backpressure
- Licensing throughput limits
- Platform limits

Refer to the section Quantum Flow processor (QFP) Drops from this document to further troubleshoot this condition.

Quantum Flow Processor (QFP) Drops

To verify reasons for QFP drops use the command **show platform hardware qfp active statistics drop** like demonstrated here:

```
---- show platform hardware qfp active statistics drop  ----

Last clearing of QFP drops statistics : never

-------------------------------------------------------------------------
Global Drop Stats                          Packets                  Octets
-------------------------------------------------------------------------
BFDoffload                                23944858              1904416850
IpTtlExceeded                               184211                28644972
IpsecIkeIndicate                               175                   26744
IpsecInput                                  686112               171458640
IpsecInvalidSa                                   1                      80
Ipv4Martian                                      4                     392
Ipv4NoAdj                                    19776                 6587643
Ipv4NoRoute                                     75                   10950
Ipv6NoRoute                                  27068                 1515808
ReassDrop                                  3489529                450382594
ReassNoFragInfo                            4561070               6387610348
ReassOverlap                                     3                     198
ReassTimeout                               7408271               2631950860
TailDrop                                 193769387             157113756882
```

This command shows different reasons for QFP drops and the associated packet counters for each category.

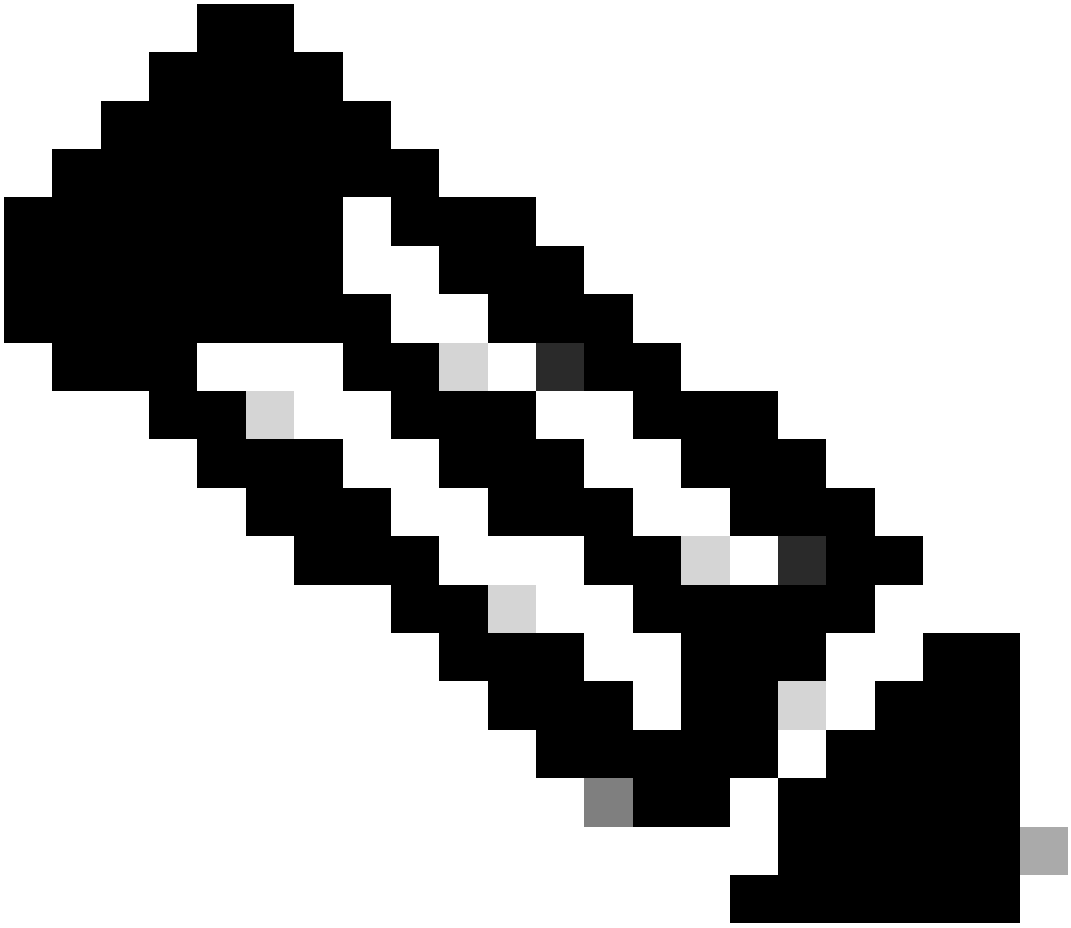

Note: Most of the QFP drop category reasons are self explanatory by its name. The reason category guides the troubleshooting flow. For non common packet drop categories, if required, file a Cisco TAC case.

Tail Drops

One of the most frequently observed drop types is the TailDrop counter, which usually increases due to these reasons:

- Licensing throughput limitation.
- QoS drops

Type of log generated:

%BW_LICENSE-4-THROUGHPUT_MAX_LEVEL: F0/0: cpp_ha_top_level_server: Average throughput rate approached the licensed bandwidth of <mbps> Mbps during <sampling-number> sampling periods in the last <period> hours, sampling period is <sampling-period> seconds

Verification commands:

- **show platform hardware qfp active infrastructure bqs queue output default interface <interface>**
- **show plat hardware qfp active infrastructure bqs queue output default all**
- **show platform hardware qfp active feature lic-bw oversubscription**
- **show platform hardware throughput level**
- **show platform hardware throughput crypto**

- Overutlization (Platform Limitation)

To understand if the impacted traffic is being dropped and a more detail packet handling view by the QFP , you can use packet trace feature. Refer to Troubleshoot with the Cisco IOS-XE Datapath Packet Trace Feature.

Type of log generated:

%IOSXE_QFP-2-LOAD_EXCEED: Slot: 0, QFP:0, Load <load-percentage>% exceeds the setting threshold.

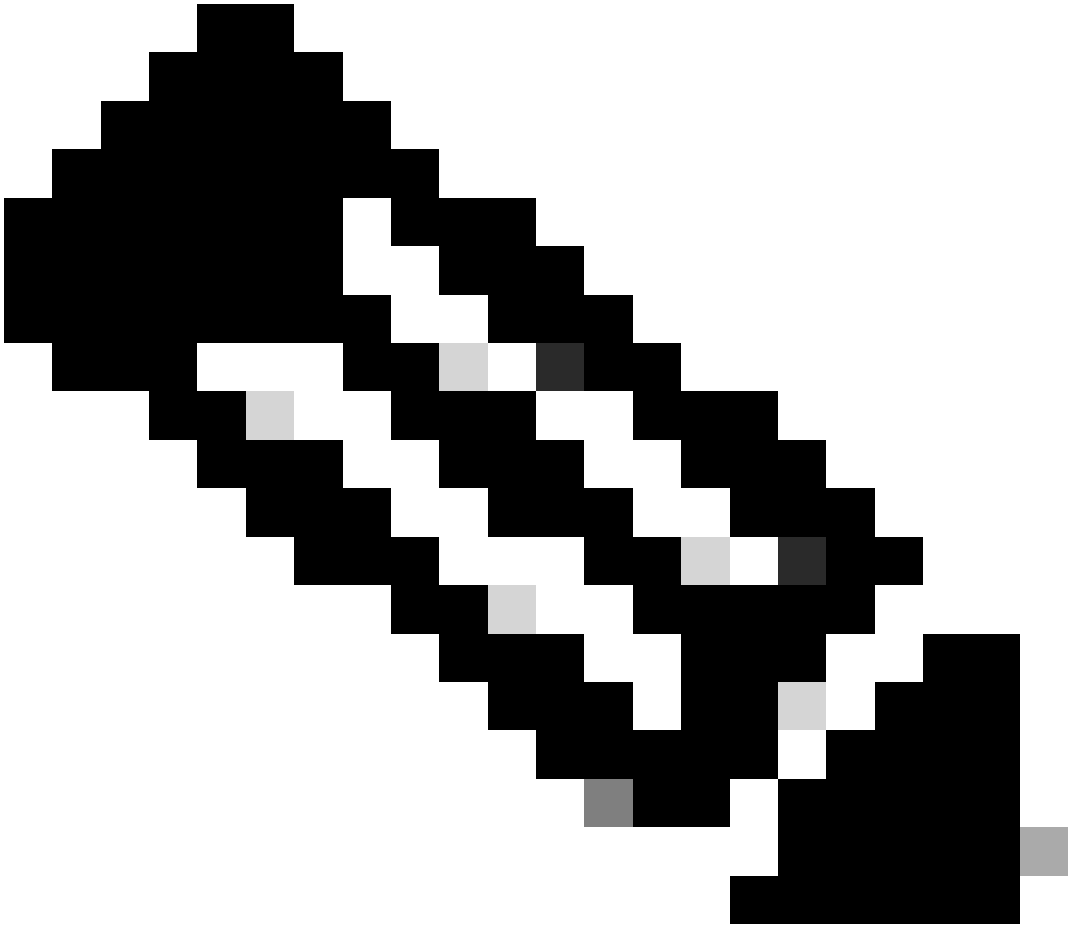

Note: Refer to platform limitation throughput numbers and scaling in datasheets. The throughput varies depending on the number and usage of features in the configuration of the device. It also vary depending on the aggregated bits per second (bps) injected into the QFP.

You can use the command **show platform hardware qfp active datapath utilization summary** the verify the QFP utilizaiton in the last 5 seconds, 1 minute, 5 minutes or 60 minutes.

```
---- show platform hardware qfp active datapath utilization summary  ----

  CPP 0:                     5 secs        1 min        5 min       60 min
Input:     Total (pps)            1            2            2            2
                 (bps)          320         1032         1032         1032
Output:    Total (pps)            0            1            1            1
                 (bps)            0         8560         8560         8576
Processing: Load (pct)            0            0            0            0

Crypto/IO
    Crypto: Load (pct)            0            0            0            0
        RX: Load (pct)            0            0            0            0
        TX: Load (pct)            2            2            2            2
            Idle (pct)           97           97           97           97
```

For additional drops verification in QFP, use **show drops { bqs | crypto| firewall| interface| ip-all| nat| punt| qfp| qos|history}** command, refer to the guide Cisco Catalyst 8500 and 8500L Series Edge Platforms Software Configuration Guide.

Additional Tips for Troubleshooting Packet Drops

Counters on Interfaces

Different counters via **show interfaces [interface]** command are shown. The explanation on the meaning on each of those counters can be found at Troubleshooting Ethernet document.

History Bits per Second CLI-based Graph

You can enable history bits per second graph view within the CLI in the incoming and outgoing direction from an interface using the command **history bps** under interface level. This configuration generates a graph of historical bps on the interface.

```
Router(config)#interface gigabitEthernet 0/0/0
Router(config-if)#history bps
```

You can also enable **history bps output-drops** and other counters history.

To display history counter results over time, use the command **show interfaces <interface> history** command:

```
---- show interfaces gigabitEthernet 0/0/0 history ?  ----

  60min   Display 60 minute histograms only
  60sec   Display 60 second histograms only
  72hour  Display 72 hour histograms only
```

```
  all     Display all three histogram intervals
  both    Display both input and output histograms
  input   Display input histograms only
  output  Display output histograms only
  |       Output modifiers
  <cr>    <cr>


#show int gi1 history 60sec
<snip>

90100                                                          *
82100                                                          *
74100                                                        ******
66100                                                      ***********
58100                                                      ***********
50100                                                    ***************
42100                                                  *********************
34100                                 ********************************
26100                              ************************************
18100          ************************************************
10100 ************************************************************
     0....5....1....1....2....2....3....3....4....4....5....5....6
               0    5    0    5    0    5    0    5    0    5    0
<snip>
      GigabitEthernet1 output rate(mbits/sec) (last 60 seconds)
```

Clear Counters

These commands are used to clear various counter statistics:

- **clear counters**: Clears interface-level counter data.
- **show platform hardware qfp active statistics drop clear**: Resets drop counters on the QFP (Quantum Flow Processor).

Percentage of Drops over Time

To verify if the number of output drops is impactful , besides using the **history bps output-drops** command at interface level to understand drops over time, you can use other counter values to get the overall percentage of output drops since the last time counters were cleared.

If the counters have not been cleared since the last boot, use the command **show version** to get the uptime of the system or refer to the last clearing of show interface counters value with the command **show interfaces.**

After determining when the counters were last cleared or identifying device uptime, calculate the percentage of output drops that occurred during that period.

This can be done by multiplying the total output drops value by 100 and then dividing the result between the packets output counter value from **show interfaces <interface>** command. The result of this operation gives an idea about the % of output drops for that interface during that time frame.

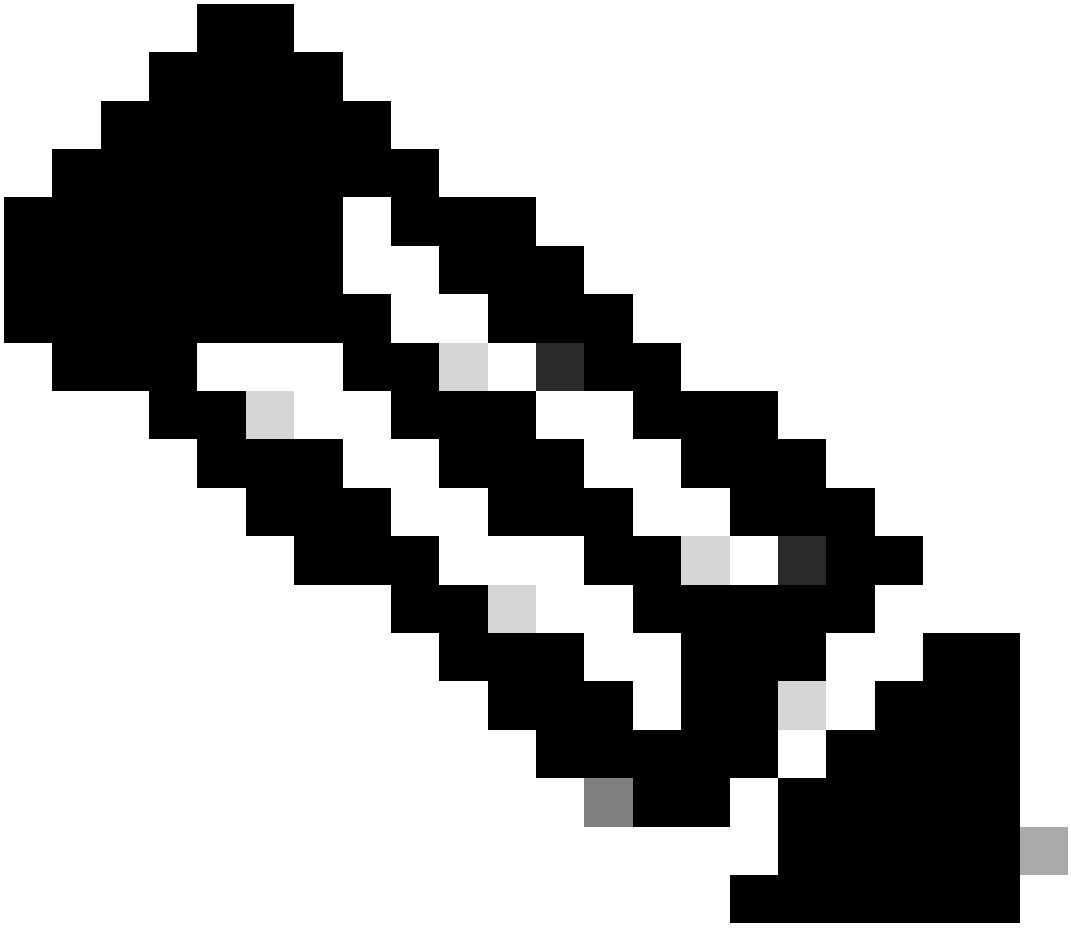

Note: Keep in mind that counters from **show interfaces** and **show platform hardware qfp active statistics drop** are historical and cumulative since the last time they were cleared. Counters are cleared if a reload is done.

Refer to this example output:

```
---- show version  ----

<snip>
Hostname uptime is 51 weeks, 1 day, 14 hours, 17 minutes
<snip>

---- show interface GigabitEthernet0/0/1  ----

GigabitEthernet0/0/1 is up, line protocol is up
<snip>
  Last clearing of "show interface" counters never
  Input queue: 0/375/0/0 (size/max/drops/flushes); Total output drops: 1351
<snip>
     219128599 packets output, 84085726336 bytes, 0 underruns
```

The example output indicates that the counters of the interfaces have never been cleared, meaning that for the last 51 weeks of uptime of the device, the percentage of total output drops is (1351 x 100) / 219128599 = .0006%.

The interpretation of this percentage can be that the total output drops on this interface is not significant and since this counter is historical, cumulative, and given the prolonged uptime, this means that the drops are non impactful.

Load Interval

Load interval is a configuration parameter from interface level which indicates the length of time for which data is used to compute load statistics.

```
Router#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Router(config)#interface gigabitEthernet 0/0/0
Router(config-if)#load-interval ?
  <30-600>  Load interval delay in seconds
```

The result of the load-interval parameter is reflected with the command **show interfaces** under the input and output rate values:

```
---- show interfaces gigabitEthernet 0/0/0  ----

GigabitEthernet0/0/0 is administratively down, line protocol is down
<snip>
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
```

This is important at the moment of verifying the bits per second rate when executing the **show interfaces** command.

The input and output rate values are helpful to understand the bits per second incoming and outgoing from an interface.

Use the command **show interface summary** for a broad overview for input and output rate values on all interfaces and get the aggregated output rate from physical interfaces which is helpful to understand total aggregated bits per second output at certain point in time. Refer to the RXBS and TXBS counters from this example output:

```
---- show interfaces summary  ----

 *: interface is up
 IHQ: pkts in input hold queue     IQD: pkts dropped from input queue
 OHQ: pkts in output hold queue    OQD: pkts dropped from output queue
 RXBS: rx rate (bits/sec)          RXPS: rx rate (pkts/sec)
 TXBS: tx rate (bits/sec)          TXPS: tx rate (pkts/sec)
```

```
 TRTL: throttle count

  Interface                  IHQ         IQD         OHQ         OQD       RXBS       RXPS       TXBS       TXPS
-----------------------------------------------------------------------------------------------------------------
* GigabitEthernet0/0/0         1           0           0           0       9000         19          0          0
  GigabitEthernet0/0/1         0           0           0           0          0          0          0          0
  GigabitEthernet0/0/2         0           0           0           0          0          0          0          0
* GigabitEthernet0/0/3         0           0           0           0       9000         19          0          0
```

Remove the QoS Policy Temporarily

To troubleshoot, temporarily remove the QoS policy from the affected interface. Use the command **no service-policy output <policy-name>** at the interface configuration level.

- If drops persist without the QoS policy, the issue is unrelated to QoS.
- If drops cease after removing the QoS policy, the policy is the likely cause.

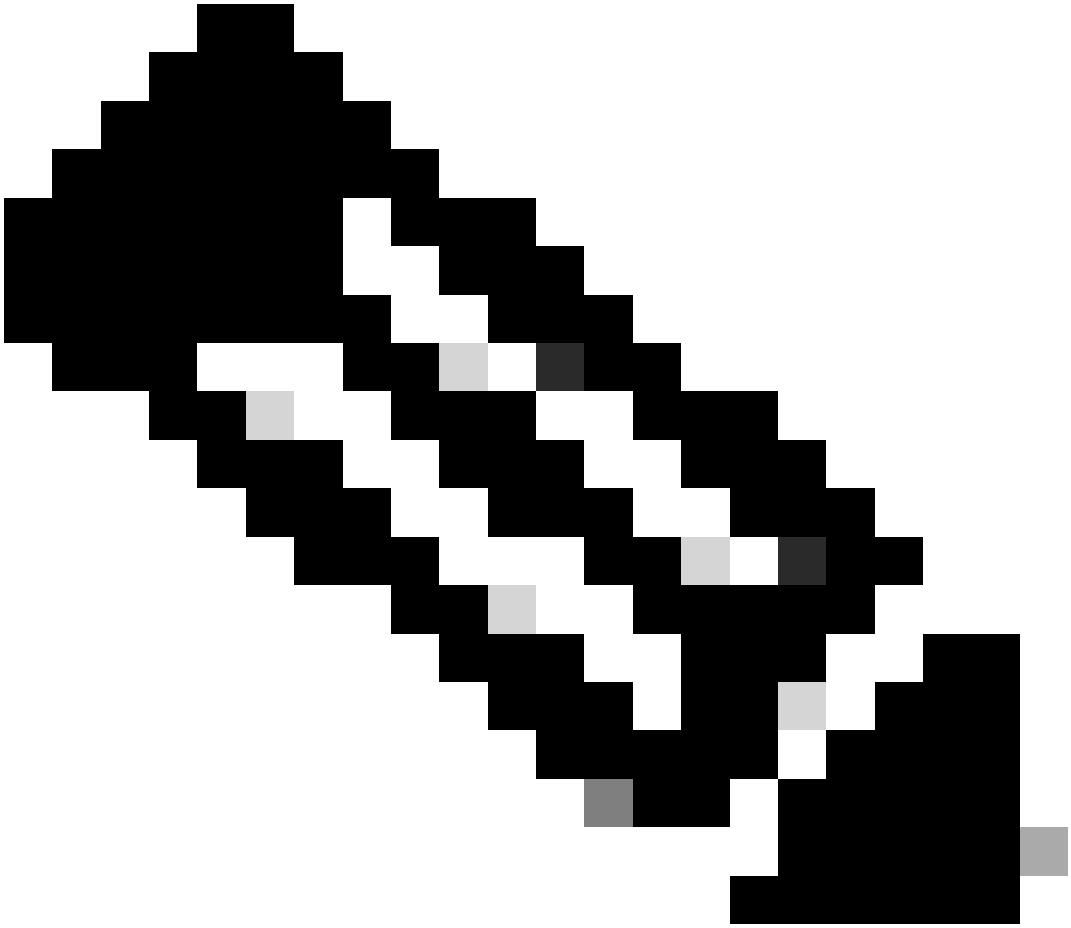

Note: If TAC assistance is required, the isolation of determining if drops are due to quality of service or not is important to route the case to the appropriate expert at early stages.

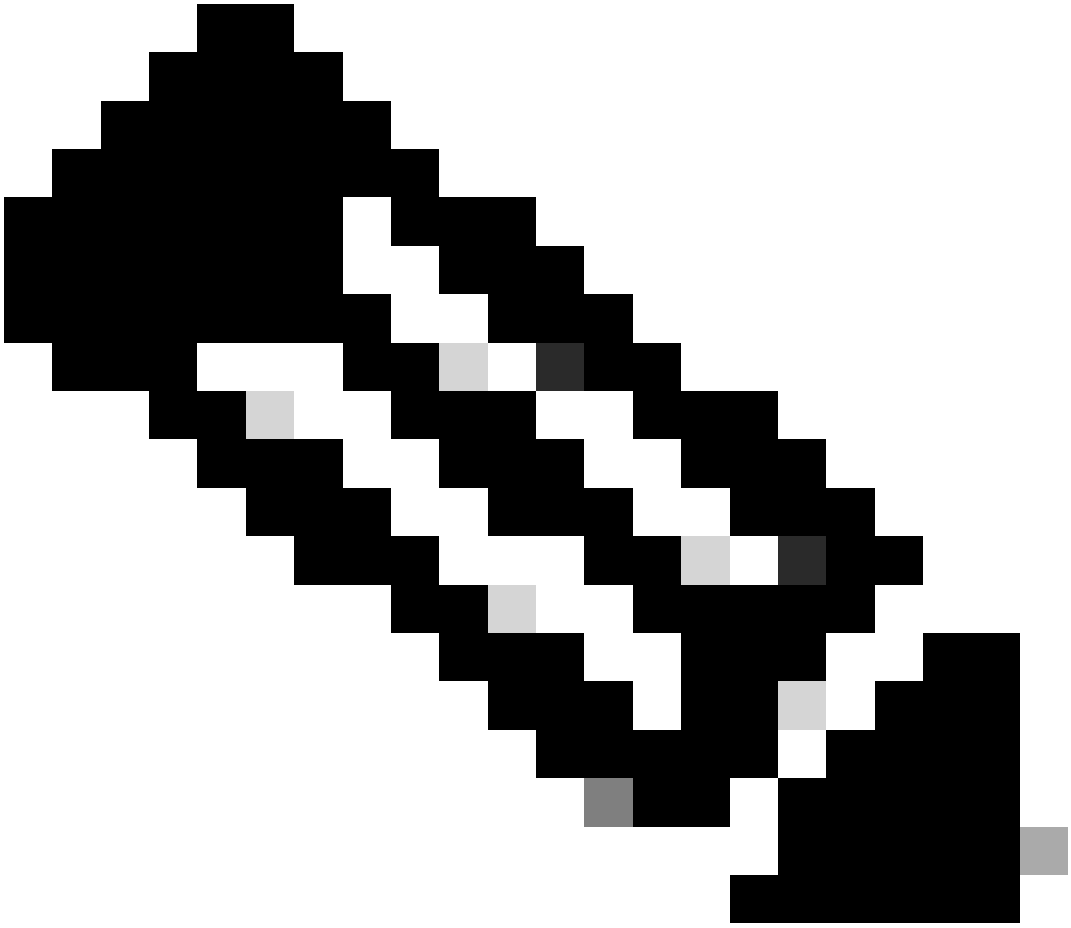

Note: Drops can also be due to ipsec feature. Ipsec drops are generally aggregated in physical interface that is used as tunnel source. If the drops are present only when the tunnel is used, this is important to indicate to TAC if assistance is required. This helps to route the case to the corresponding team at early stages

Related Articles and Documentation

- Troubleshoot Packet Drops on ASR 1000 Series Service Routers