Troubleshoot High-Utilization Due to NAT Gatekeeper Default Configuration

Contents

Introduction

Prerequisites

- Requirements
- Components Used

Background Information

Symptoms

Packet Trace Feature

- Basic Packet Trace Configuration

What is the NAT Gatekeeper

Check the NAT Gatekeeper

Workaround/Fix

- Solution 1
- Solution 2

Summary

Related Information

Introduction

This document describes how to identify and resolve High QFP utilization on Routing Platforms caused by mix of NATed and non-NATed traffic.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Basic knowledge of Cisco IOS® XE packet forwarding architecture
- Basic experience with Packet Trace Feature

Components Used

This document is not restricted to specific software and hardware versions. It applies for any routing Cisco IOS XE platform with physical/virtualized High Quantum Flow Processor (QFP) like ASR1000, ISR4000, ISR1000, Cat8000 or Cat8000v.

This document is based on Cisco IOS XE devices in autonomous mode, SDWAN (controller) or SD-routing can adhere to a similar logic but specifics can be different.

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

High-utilization and performance issues on the Cisco Quantum Flow Processor (QFP) can be observed on a Cisco router when there is a mix of NATed and Non-NAT traffic flows present on the same interface. This can also lead to other performance issues such as interface errors or slowness.



Note: The QFP is located on the Embedded Services Processor (ESP) and it is in charge of the data plane and packet processing for all the inbound and outbound traffic flows, this can be either physical or virtualized depending on the platform.

Symptoms

It is important to validate and confirm these symptoms from the router in order to identify this behavior:

1.- HIgh QFP Load alerts. These alerts appear when the Load exceeds the threshold of 80%.

```
Feb 8 08:02:25.147 mst: %IOSXE_QFP-2-LOAD_EXCEED: Slot: 0, QFP:0, Load 81% exceeds the setting threshol
Feb 8 08:04:15.149 mst: %IOSXE_QFP-2-LOAD_RECOVER: Slot: 0, QFP:0, Load 59% recovered.
```



Note: You can also run the **show platform hardware qfp active datapath utilization summary** command in order to reveal the load on the QFP and the traffic rates.

```
Router# show platform hardware qfp active datapath utilization summary
  CPP 0: Subdev 0            5 secs        1 min        5 min       60 min
Input:  Priority (pps)            0            0            0            0
                 (bps)           96           32           32           32
    Non-Priority (pps)       327503       526605       552898       594269
                 (bps)   1225600520   2664222472   2867573720   2960588728
           Total (pps)       327503       526605       552898       594269
                 (bps)   1225600616   2664222504   2867573752   2960588760
```

```
Output: Priority (pps)            6            7            7            7
                 (bps)         8576         9992         9320         9344
    Non-Priority (pps)       327715       526839       553128       594506
                 (bps)   1257522072   2714335584   2920005904   3016943800
           Total (pps)       327721       526846       553135       594513
                 (bps)   1257530648   2714345576   2920015224   3016953144
Processing: Load (pct)           99           72           34           19
```

2.- Interface errors. Packets can be dropped due to backpressure If there is high QFP utilization. In such cases, Overruns and Input Drops are commonly observed on the interfaces. To display this information, you can run the **show interfaces** command.

```
Router# show interface gigabitEthernet 0/0/1
GigabitEthernet0/0/1 is up, line protocol is up
  Hardware is ISR4351-3x1GE, address is e41f.7b59.cba1 (bia e41f.7b59.cba1)
  Description: ### LAN Interface ###
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 2/255
  Encapsulation 802.1Q Virtual LAN, Vlan ID  1., loopback not set
  Keepalive not supported
  Full Duplex, 1000Mbps, link type is force-up, media type is LX
  output flow-control is on, input flow-control is on
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:02, output 00:06:47, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/375/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  30 second input rate 9390000 bits/sec, 2551 packets/sec
  30 second output rate 1402000 bits/sec, 1323 packets/sec
     368345166434 packets input, 199203081647360 bytes, 0 no buffer
     Received 159964 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     2884115457 input errors, 0 CRC, 0 frame, 2884115457 overrun, 0 ignored
     0 watchdog, 3691484 multicast, 0 pause input
     220286824008 packets output, 32398293188401 bytes, 0 underruns
     0 output errors, 0 collisions, 4 interface resets
     3682606 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     21 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapped out
```

3.- In some scenarios, users can complain of slowness on the network.

Packet Trace Feature

Packet Trace is a tool that provides detailed information of how data packets are processed by the Cisco IOS XE platforms. It has 3 levels of inspection which are accounting, summary, and data path. The level of inspection is based on the debug platform condition state.

You can obtain information like:

- Input and output interface
- Packet state
- Timestamps
- Packet trace

Note: Configure data path consumes more packet-processing resources, which is only reflected on the packets that match the filter condition.

More details about Packet Trace in Troubleshoot with the Cisco IOS XE Datapath Packet Trace Feature.

Basic Packet Trace Configuration

This is an example of a basic Packet Trace configuration with Data Path level inspection. It collects 8192 packets in a circular way (overwrites old packets), creates a copy of each packet from Layer 3 that arrives and leaves interface GigabitEthernet 0/0/1.

```
Router# debug platform packet-trace packet 8192 circular fia-trace data-size 2048
Router# debug platform packet-trace copy packet both L3 size 64
Router# debug platform condition interface gigabitEthernet 0/0/1 both
Router# debug platform condition start
Router# debug platform condition stop
```

You can check the results of Packet Trace with these commands.

```
Router# show platform packet-trace summary
Router# show platform packet-trace packet all
```

From Packet Trace capture, you can observe that the NAT feature consumes more resources than expected. In the next example, you can see that the lapsed time for the IPV4_NAT_INPUT_FIA feature is significantly larger than the lapsed time from other features. This behavior usually indicates that the QFP takes more time to process this feature and, as result, more resources from the QFP are used for NAT.

```
Packet: 161         CBUG ID: 161
Summary
  Input     : GigabitEthernet0/0/1
  Output    : GigabitEthernet0/0/2.1730
  State     : FWD
  Timestamp
    Start   : 25136781447706429 ns (02/10/2024 00:25:49.584050 UTC)
    Stop    : 25136781447993237 ns (02/10/2024 00:25:49.584337 UTC)
<snip>
  Feature: IPV4_NAT_INPUT_FIA   <<<<<<<<<<<<
    Entry       : Input - 0x700162ac
    Input       : GigabitEthernet0/0/1
    Output      : <unknown>
```

```
    Lapsed time : 1873376 ns     <<<<<<<<<<<<
<snip>
  Feature: IPV4_INPUT_IPOPTIONS_PROCESS
    Entry       : Input - 0x70016344
    Input       : GigabitEthernet0/0/1
    Output      : GigabitEthernet0/0/2.1730
    Lapsed time : 64 ns
```

What is the NAT Gatekeeper

In Cisco IOS XE Routing Platforms, the feature Network Address Translation (NAT) Gatekeeper is enabled by default. NAT Gatekeeper was originally created to prevent non-NAT-ed flows to use excessive Processing Resources to create a NAT translation. NAT Gatekeeper creates two small caches for the inside to outside direction and for other direction based on the source address. Each cache entry consists of a source address, a virtual routing and forwarding (VRF) ID, a timer value (used to invalidate the entry), and a frame counter.

High volume of non-NATed traffic on a NATed interface consumes a high amount of resources and causes the QFP utilization spikes. Cisco recommends that customers must not have NAT-ed and non-NAT-ed flows on the same interface wherever possible.

Check the NAT Gatekeeper

The NAT Gatekeeper Statistics can be check with the commands **show platform hardware qfp active feature nat datapath { gatein | gateout } activity**. This shows the size of the cache, the number of hits, misses, aged, added and active entries in the cache.Usually, if there is a high number of misses and if this number increases rapidly in a short period of time, this indicates that a huge number of Not-Natted flows are not added to the cache. This behavior causes these flows are processed by the QFP within the NAT workflow, and this can up in high QFP utilization.

```
Router# show platform hardware qfp active feature nat datapath gatein activity
Gatekeeper on
def mode Size 8192, Hits 191540578459, Miss 3196566091, Aged 1365537 Added 9 Active 7

Router# show platform hardware qfp active feature nat datapath gateout activity
Gatekeeper on
def mode Size 8192, Hits 448492109001, Miss 53295038401, Aged 149941327 Added 603614728 Active 1899
```

Workaround/Fix

In most environments, the NAT gatekeeper functionality works fine and does not cause issues. However, if you do run into this problem there are a few ways to resolve it.

Solution 1

For these type of issues, Cisco recommends separating the NATed and non-NATed traffic from the same interface. It can be used either in different interfaces or network devices.

Solution 2

Increase the size of the cache on the NAT Gatekeeper feature in order to reduce the number of misses from the gatekeeper.

The next example shows how to adjust the Gatekeeper on a Cisco router. Please note this value must be represented in powers of 2. Otherwise, the value automatically sets to the next lower size.

```
Router(config)# ip nat service gatekeeper
Router(config)# ip nat settings gatekeeper-size 65536
```



Note: Adjusting the cache size can cost exmem memory within the QFP, so optimize its usage. Try to adjust this value gradually and start with the nearest possible value to the default setting.

After performing one of the solutions described, it is recommended to monitor these two parameters in order to confirm the issue has been resolved:

1.- Verify that the QFP utilization has decreased.

2.- Verify that the number of *misses* are not continue to increase.

Summary

The NAT Gatekeeper feature can enhance performance of the router when there is non-NATed flows on a NATed interface. This usually happens when NAT translates some NATed flows when, at the same time, non-NATed flows pass through the same interface. In most environments, the NAT Gatekeeper feature does not cause any impact to the router. However, it is important to adjust this feature if needed, carefully in order to avoid side effects.

Related Information

- ASR1K NAT Intermittently Fails to Translate Some Packets
- Troubleshoot with the Cisco IOS XE Datapath Packet Trace Feature
- Cisco Technical Support & Downloads