

Configure QoS Priority in ASR 920 and RSP2 Based Routers

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Problems](#)

[Priority Packets Dropped at The Advantage of Best Effort Packets](#)

[Solution](#)

[How to Calculate The Bandwidth Percent to Offer in a Class](#)

[RSP2 QoS Priority Test](#)

[Configuration](#)

[Traffic](#)

[Frame Example](#)

[Test with Traffic Generator](#)

[Negative Test](#)

[Conclusons](#)

Introduction

This document describes how ASR920/RSP2 routers handle QoS priorities and how to configure it.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- ASR 920 series routers
- QoS policies

Components Used

The information in this document is based on a ASR 9xx with RSP2 router that runs software version 16.x through 17.x.

A traffic generator is used to test the functions that handle priority packets.

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.

Problems

This document addresses these problems specific of ASR 920 and RSP2 based routers:

- Priority packets dropped at the advantage of a best-effort packets due to ASIC limitation on RSP2
- How to calculate the bandwidth percent to offer in a class

Priority Packets Dropped at The Advantage of Best Effort Packets

During a test it was determined that priority packet can be dropped at the advantage of a best-effort packet. This is apparent when incoming traffic arrives through an interface with higher speed than the egress interface and cause over-subscription in the output direction. For example, when 5 Gbps of traffic is received and needs to be forwarded through a 1 Gbps interface.

This is equally true for egress interfaces configured with a shaper. In case the ingress speed is higher than the configured CIR at the egress priority, a packet could still be dropped at the advantage of a best effort packet.

Note: There is an ASIC limitation for which we cannot have a priority child to a non-priority parent.

If a queue is configured as expedite and its parent sub-channel is not, there is jitter on the priority queue due to the arbitration latencies at the sub-channel level.

Solution

1. Configure EFP
2. Apply a shaper on the physical
3. Apply the desired QoS on the EFP
4. Apply the IP connectivity in the BDI interface

Example:

```
configuration with issue
```

```
-----  
interface GigabitEthernet0/0/0  
description this is my egress interface  
service-policy output PM-1G-Out
```

```
configuration without issue
```

```
-----  
interface GigabitEthernet0/0/0  
description this is egress interface  
service-policy output POL-PRIO-MAIN-1G ==> shaper, useful to allow internal priority like BDF  
service instance 200 ethernet  
encapsulation dot1q 200  
rewrite ingress tag pop 1 symmetric  
service-policy output PM-1G-Out ==> the original QoS previously applied in the physical  
interface  
bridge-domain 200
```

```
!  
interface BDI200 ==> BDI must match the bridge-domain defined under the service-instance  
description this is L3 egress  
ip address 10.20.2.45 255.255.255.0  
ip mtu 9000  
=> no QoS applied under the BDI, all QoS are in the service-instance with a backpressure of the  
shaper in the physical
```

With this configuration, all priority packets were prioritized properly and none was dropped at the advantage of a best effort packet, still you need to calculate the bandwidth allocated properly.

How to Calculate The Bandwidth Percent to Offer in a Class

The bandwidth allocation in RSP2 platform has also a specific behavior. Many times drops are seen while the QoS is configured as in other platforms.

For example, if you configure QoS with a shaper of 2 Mbps in a ASR1K router, it does not drop before 2 Mbps are reached, nor it queues packets in the class. However, this happens with RSP2.

In many cases, the offered speed does not even reach the maximum allowed when drops are already seen.

This is a typical example of what can be seen on a RSP2, while the same values for the exact same traffic applied to another platform would not show any drop:

```
ASR903#show ethernet service instance policy-map | s EXP-5  
Class-map: EXP-5 (match-all)  
58803127 packets, 5488269944 bytes  
5 minute offered rate 279000 bps, drop rate 35000 bps  
  
Match: mpls experimental topmost 5  
Priority: 3% (297 kbps), burst bytes 37000, b/w exceed drops: 60373  
  
Priority Level: 1
```

The issue is due to the way traffic is handled in the hardware. Basically, RSP2 hardware implementation does not consider only the layer 3 but the entire frame, which means all headers are taken into account.

RSP2 QoS Priority Test

In this case CEM traffic is used to test the priority behavior.

This is an example that shows how to configure priority to avoid drops at the advantage of best-effort and tune the bandwidth allocation.

Configuration

```
policy-map POL-PRIO-MAIN-1G  
class class-default  
shape average 8650000  
!  
policy-map PM-MPLS-1G-Out  
class EXP-5  
priority level 1 percent 4
```


```

class EXP-4
priority level 2 percent 24
class EXP-6
bandwidth percent 2
queue-limit 25000 us
class EXP-3
bandwidth percent 2
queue-limit 10000 us
class EXP-2
bandwidth percent 2
queue-limit 50000 us
class EXP-1
bandwidth percent 2
queue-limit 20000 us
class class-default
bandwidth percent 1
queue-limit 40000 us
!
interface GigabitEthernet0/0/0
no ip address
negotiation auto
service-policy output POL-PRIO-MAIN-1G
service instance 200 ethernet
encapsulation dot1q 200
rewrite ingress tag pop 1 symmetric
service-policy output PM-MPLS-1G-Out
bridge-domain 200
!
interface CEM0/1/8
no ip address
cem 0
service-policy input PM-CEM-in
payload-size 128
dejitter-buffer 20
!
interface CEM0/1/9
no ip address
cem 0
service-policy input PM-CEM-in
payload-size 64
dejitter-buffer 16
!
interface BDI200
description path for qos stress
ip address 10.20.2.45 255.255.255.0
ip mtu 9000
ip router isis
carrier-delay msec 0
cdp enable
mpls traffic-eng tunnels
bfd template BFD-1hop-5ms
isis circuit-type level-2-only
isis network point-to-point
isis metric 15 level-1
isis metric 15 level-2
ip rsvp bandwidth percent 90
ip rsvp signalling hello graceful-restart

```

Traffic

2 streams of traffic are created by CEM0/1/8 in red and CEM0/1/9 in green:



We can see the behavior with different packet size, CEM0/1/9 sends twice as much packets as CEM0/1/8, which is configured for 128 bytes.

Normally, a QoS configuration on RP considers only the payload of CEM, RSP2 considers the entire frame instead.

Frame Example

The left screenshot shows a packet with the following details:

- Frame 8: 94 bytes on wire (752 bits), 94 bytes captured (752 bits) on interface ...
- Ethernet II, Src: Cisco_3b:f2:bf (70:1f:53:3b:f2:bf), Dst: Cisco_a6:04:bf (f8:7b:20:a6:04:bf)
- Destination: Cisco_a6:04:bf (f8:7b:20:a6:04:bf)
- Source: Cisco_3b:f2:bf (70:1f:53:3b:f2:bf)
- Type: 802.1Q Virtual LAN (0x8100)
- 802.1Q Virtual LAN, PRI: 5, DEI: 0, ID: 200
- MultiProtocol Label Switching Header, Label: 6017, Exp: 5, S: 1, TTL: 255
- PW Ethernet Control Word
 - Sequence Number: 50785
 - Ethernet II, Src: e7:e7:e7:e7:e7:e7 (e7:e7:e7:e7:e7:e7), Dst: e7:e7:e7:e7:e7:e7 (e7:e7:e7:e7:e7:e7)
 - Destination: e7:e7:e7:e7:e7:e7 (e7:e7:e7:e7:e7:e7)
 - Source: e7:e7:e7:e7:e7:e7 (e7:e7:e7:e7:e7:e7)
 - Type: Unknown (0xe7e7)
 - Data (54 bytes)

The right screenshot shows a packet with the following details:

- Frame 6: 158 bytes on wire (1264 bits), 158 bytes captured (1264 bits) on interface ...
- Ethernet II, Src: Cisco_3b:f2:bf (70:1f:53:3b:f2:bf), Dst: Cisco_a6:04:bf (f8:7b:20:a6:04:bf)
- Destination: Cisco_a6:04:bf (f8:7b:20:a6:04:bf)
- Source: Cisco_3b:f2:bf (70:1f:53:3b:f2:bf)
- Type: 802.1Q Virtual LAN (0x8100)
- 802.1Q Virtual LAN, PRI: 5, DEI: 0, ID: 200
- MultiProtocol Label Switching Header, Label: 6001, Exp: 5, S: 1, TTL: 255
- PW Ethernet Control Word
 - Sequence Number: 31601
 - Ethernet II, Src: fc:fc:fc:fc:fc:fc (fc:fc:fc:fc:fc:fc), Dst: fc:fc:fc:fc:fc:fc (fc:fc:fc:fc:fc:fc)
 - Destination: fc:fc:fc:fc:fc:fc (fc:fc:fc:fc:fc:fc)
 - Source: fc:fc:fc:fc:fc:fc (fc:fc:fc:fc:fc:fc)
 - Type: Unknown (0xfcfc)
 - MDS Header(Unknown(12)/Unknown(11))
 - Fibre Channel

You can see 30 bytes extra to the original payload configured under CEM. This can be explained as:

Ethernet header = 14 Bytes
 Dot1q header = 4 Bytes

Mpls header = 4 Bytes
PW Header = 4 Bytes
CEM trailer = 4 Bytes
Total = 30 Bytes

Calculation of the bandwidth needed in the hardware, remind that the frame needs to be considered:

CEM 0/1/8 125 Packet/sec, size 128bytes ==> $125 * 128 * 8 = 128000$ bps
CEM 0/1/9 250 Packet/sec, size 64bytes ==> $250 * 64 * 8 = 128000$ bps
on each frame we need an extra 30bytes ==> $375 * 30 * 8 = 90000$ bps
Total = 346000 bps

To verify the behavior and accuracy of the shaper on the interface it was configured to 8650000 bps, this is done to have an exact 4% for priority class.

Calculation: $346000.0000 / 8650000.0000 = 0.04 = 4\%$.

This is what is seen in the configuration above. The results confirm that the configuration and calculation are accurate.

Policy output:

```
ASR903#show ethernet service instance policy-map | s EXP-5
Class-map: EXP-5 (match-all)
3063745 packets, 285949512 bytes
5 minute offered rate 279000 bps, drop rate 0000 bps
Match: mpls experimental topmost 5
Priority: 4% (346 kbps), burst bytes 8650, b/w exceed drops: 0
```

Priority Level: 1

346 Kbps applied in platform independent is way more than L3 but is exactly the L2 traffic.

Test with Traffic Generator

Traffic generator --> TenGig interface --> Asr9xx RSP2 --> output 1G where the policy is applied.

```
ASR903#show clock
22:54:40.976 CET Wed Nov 30 2022
ASR903#show ethernet service instance policy-map | inc Class-map:|drop rate
Class-map: EXP-5 (match-all)
5 minute offered rate 279000 bps, drop rate 0000 bps
Class-map: EXP-4 (match-all)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: EXP-6 (match-any)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: EXP-3 (match-any)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: EXP-2 (match-all)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: EXP-1 (match-any)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: class-default (match-any)
5 minute offered rate 762348000 bps, drop rate 756024000 bps
```

```
ASR903#show clock
17:41:16.110 CET Thu Dec 1 2022
```

```
ASR903#show ethernet service instance policy-map | inc Class-map:|drop rate
Class-map: EXP-5 (match-all)
5 minute offered rate 279000 bps, drop rate 0000 bps
Class-map: EXP-4 (match-all)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: EXP-6 (match-any)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: EXP-3 (match-any)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: EXP-2 (match-all)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: EXP-1 (match-any)
5 minute offered rate 0000 bps, drop rate 0000 bps
Class-map: class-default (match-any)
5 minute offered rate 762400000 bps, drop rate 756077000 bps
```

After approximately 18 hours there was not a single drop in priority, although on the interface there are a lot of drops, as seen in the drop rate of the class-default, due to the CIR of the shaper limit.

Notice that the default queue-limit was used: in order to tune the bandwidth to support the entire L2 frame size, you do not require to tune the queues.

Negative Test

Another test to verify the good accuracy is to omit the 4 bytes of the CEM trailer and see if small drops occur:

```
ASR903#show ethernet service instance policy-map | s EXP-5
Class-map: EXP-5 (match-all)
352466 packets, 32896848 bytes
5 minute offered rate 279000 bps, drop rate 5000 bps
Match: mpls experimental topmost 5
Priority: 4% (334 kbps), burst bytes 8350, b/w exceed drops: 271

Priority Level: 1
```

As you can see, if you omit a part of that frame it causes drops.

Conclusions

This test with the CEM traffic confirms that the entire L2 frame needs to be taken in consideration for the bandwidth calculation.

One artifice is to increase the queue-limit, but clearly a correct calculation of the L2 frame puts less stress on the memory resources used by the platform.

It is obvious that not all traffic can be foreseen at every point in time, like in transfer with variable packet size. In order to have an accurate configuration, you need to take into consideration ethernet, dot1q(s), MPLS tag(s) headers to the average packet size and also the packets rate.

For any traffic that traverses the ASIC of an RSP2, you need to be aware of every single byte that is contained in a frame sent out of the platform (CRC not included).