

Monitoring Queue Thresholds

Purpose Display the color-based thresholds for queues on each egress slot.

Showing queue thresholds by queue profile shows buffer memory information for each queue profile and, within that profile, shows the thresholds for each region.

In addition, showing queue thresholds by region organizes the buffer memory information by queue region and, within each region, shows the buffer allocations for each queue profile.

Action To display the color-based queue thresholds for each of the 2000 video queues when 8000 total queues are configured:

```
host1#show qos queue-thresholds egress-slot 9 queue-profile video
queue-profile video 2000 queues
```

region	egress memory	exceeded length	conformed length	committed length	total committed memory
0	0MB - 4MB	34944	69888	139648	279296000
1	4MB - 8MB	24448	48896	97792	195584000
2	8MB - 12MB	14080	28032	55936	111872000
3	12MB - 16MB	7040	14080	28032	56064000
4	16MB - 20MB	5248	10496	20992	41984000
5	20MB - 24MB	1280	2560	5120	10240000
6	24MB - 28MB	1152	2176	4224	8448000
7	28MB - 32MB	896	1792	3456	6912000

As shown, when all of the egress memory in use is between 0 MB and 4 MB, each video queue can queue 139,648 bytes of committed traffic. Because the default conformed fraction is 50 percent and the default exceeded fraction is 25 percent, half of the committed length, or 69,888 bytes, can be queued before conformed traffic is dropped, and one quarter of the committed length, or 34,944 bytes, can be queued before exceeded traffic is dropped. While memory fills, the video queues are given progressively smaller amounts of memory. For example, when 28 to 32 MB of buffer memory is in use, each video queue is limited to 3456 bytes. While memory fills beyond the last region, all frames are dropped except control traffic, until the queues are drained and memory usage falls back into one of the regions.

To display the router's memory management:

```
host1#show qos queue-thresholds egress-slot 9 region 0
region 0 (0MB - 4MB) oversubscription 3330%
```

queue-profile	exceeded length	conformed length	committed length	queue count	total committed memory
default	34944	69888	139648	2000	279296000
video	34944	69888	139648	2000	279296000
multicast	34944	69888	139648	2000	279296000
internet	34944	69888	139648	2000	279296000

Static and dynamic oversubscription determines that when 8000 queues are configured and 0–4 MB of egress buffer memory is in use, memory is oversubscribed by 3330 percent. If significantly fewer queues are configured, there is less oversubscription. This example illustrates static oversubscription.

Because all of the queues in Example 2 use default queue profiles, all queues have the same lengths. Each queue is allocated 139,648 bytes of committed buffer memory when operating within this region. This allocation allows active queues to burst traffic by using memory that is unused by quiescent queues. This example illustrates dynamic oversubscription, which is based on the assumption that when a large number of queues is configured, only a fraction of the queues is active at a given time. While more queues become active, memory fills and spills into another region. When this occurs, queues are given progressively smaller queue limits.

In memory regions 1 through 5, queue limits are progressively reduced. In region 6, memory is strictly partitioned among queues.

To display oversubscription in region 6:

```
host1#show qos queue-thresholds egress-slot 9 region 6
region 6 (24MB - 28MB) oversubscription 100%
```

queue-profile	exceeded length	conformed length	committed length	queue count	total committed memory
default	1152	2176	4224	2000	8448000
video	1152	2176	4224	2000	8448000
multicast	1152	2176	4224	2000	8448000
internet	1152	2176	4224	2000	8448000

Oversubscription is 100 percent. When 24–28 MB of the memory is in use, there is no oversubscription of egress buffer memory; 32 MB of the 32-MB memory is allocated. In Example 3, each of the 8000 egress queues is given a queue of 4224 bytes, for a total of 16 MB.

If memory continues to fill into region 7, egress buffer memory is undersubscribed, allowing control traffic to flow within the router. As shown in Example 4, when operating in region 7, only 80 percent of the 32-MB memory is allocated.

To display oversubscription in region 7:

```
host1#show qos queue-thresholds egress-slot 9 region 7
region 7 (28MB - 32MB) oversubscription 80%
```

queue-profile	exceeded length	conformed length	committed length	queue count	total committed memory
default	896	1792	3456	2000	6912000
video	896	1792	3456	2000	6912000
multicast	896	1792	3456	2000	6912000
internet	896	1792	3456	2000	6912000

Region 7 has 2000 IP users, each with four queues. Each of the four queues use default queue profiles.

To display the queue thresholds in the multicast queue profile:

```
host1#show qos queue-thresholds egress-slot 9 queue-profile multicast
queue-profile multicast 2000 queues
```

region	egress memory	exceeded length	conformed length	committed length	total committed memory
0	0MB - 4MB	5120	10112	20096	40192000
1	4MB - 8MB	5120	10112	20096	40192000
2	8MB - 12MB	5120	10112	20096	40192000
3	12MB - 16MB	5120	10112	20096	40192000
4	16MB - 20MB	5120	10112	20096	40192000
5	20MB - 24MB	1280	2560	10112	20224000
6	24MB - 28MB	1152	2176	4224	8448000
7	28MB - 32MB	896	1792	3456	6912000

The multicast queue profile is configured with a committed length of 10,000 minimum and 20,000 maximum. When in regions 0–4, these queues would normally get more memory than the 20,000 byte maximum requested. In this case, the queue is limited to the maximum, and any excess memory is redistributed to other queues. Region 5 does not have enough memory to honor the 20,000-byte maximum requested.

Although a 20,000 byte maximum was requested, the router provisions memory in 128 byte blocks, rounded up or down per each request; 20,096 bytes is 157 blocks of 128 bytes.

In region 6, memory is strictly partitioned, and neither the minimum nor maximum request is honored. Instead, each multicast queue is given a fair share of the queue length so that aggressive bandwidth consumers cannot starve out moderate traffic consumers.

In region 7, memory is underprovisioned to allow queues to drain and to avoid starvation that occurs when egress buffer memory fills completely.

To display the queue thresholds for video queues:

```
host1#show qos queue-thresholds egress-slot 9 region 0
region 0 (0MB - 4MB) oversubscription 3330%
```

queue-profile	exceeded length	conformed length	committed length	queue count	total committed memory
default	33664	67328	134656	2000	269312000
video	67328	134656	269184	2000	538368000
multicast	5120	10112	20096	2000	40192000
internet	33664	67328	134656	2000	269312000

You can configure video queues with a buffer weight of 16 and Internet and multicast queues with a buffer weight of 8 to ensure that video queues get to queue twice as much traffic as Internet and multicast queues.

Meaning Table 1 lists the **show qos queue-thresholds** command output fields.

Table 1: show qos queue-thresholds Output Fields

Field Name	Field Description
queue profile	Name of the queue profile
region	Egress buffer memory region
egress memory	Amount of memory in each region
exceeded length	Amount of exceeded traffic that can be queued at this egress memory usage
conformed length	Amount of conformed traffic that can be queued at this egress memory usage
committed length	Amount of committed traffic that can be queued at this egress memory usage
total committed memory	Amount of committed memory allocated to the queue

- Related Topics**
- Configuring Queue Profiles to Manage Buffers and Thresholds
 - show qos queue-thresholds

Published: 2010-01-12