

Understanding Packet Counters in show policy-map interface Output

Document ID: 10107

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

What Is Congestion?

What Is the Difference Between "Packets" and "Packets Matched"?

How Are Conversation Numbers Allocated?

Confirming Your Service Policy

Related Information

Introduction

This document explains how to interpret the **show policy-map interface** command output, which you can use to monitor the results of a service-policy created with Cisco modular Quality of Service (QoS) command-line interface (CLI) (MQC).

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 presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before you use it.

Note: In Cisco IOS[®] Software Release 12.1T, packets in the outputs of the commands listed in this document include all packets that match a particular class. However, in Cisco IOS Software Release 12.1, only packets that are queued during congestion are counted and displayed in the output of these same commands.

Conventions

For more information on document conventions, refer to Cisco Technical Tips Conventions.

What Is Congestion?

To understand how to interpret the **show policy-map interface** command, you first need to understand the concept of congestion.

Conceptually, congestion is defined by the Cisco IOS software configuration guide as: "During periods of transmit congestion at the outgoing interface, packets arrive faster than the interface can send them."

In other words, congestion typically occurs when a fast ingress interface feeds a relatively slow egress interface. A common congestion point is a branch-office router with an Ethernet port facing the LAN and a serial port facing the WAN. Users on the LAN segment generate 10 Mbps of traffic, which is fed into a T1 with 1.5 Mbps of bandwidth.

Functionally, congestion is defined as filling the transmit ring on the interface. A *ring* is a special buffer control structure. Every interface supports a pair of rings: a receive ring for receiving packets and a transmit ring for transmitting packets. The size of the rings varies with the interface controller and with the bandwidth of the interface or virtual circuit (VC). For example, use the **show atm vc {vcd}** command to display the value of the transmit ring on a PA-A3 ATM port adapter. Refer to Understanding and Tuning the tx-ring-limit Value for more information.

```
7200-1# show atm vc 3
ATM5/0.2: VCD: 3, VPI: 2, VCI: 2
VBR-NRT, PeakRate: 30000, Average Rate: 20000, Burst Cells: 94
AAL5-LLC/SNAP, etype:0x0, Flags: 0x20, VCmode: 0x0
OAM frequency: 0 second(s)
PA TxRingLimit: 10
InARP frequency: 15 minutes(s)
Transmit priority 2
InPkts: 0, OutPkts: 0, InBytes: 0, OutBytes: 0
InPRoc: 0, OutPRoc: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
InPktDrops: 0, OutPktDrops: 0
CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0
OAM cells received: 0
OAM cells sent: 0
Status: UP
```

Cisco IOS, also referred to as the Layer 3 (L3) processor, and the interface driver use the transmit ring when moving packets to the physical media. The two processors collaborate in this way:

- The interface transmits packets in accordance with the interface rate or a shaped rate.
- The interface maintains a hardware queue or transmit ring, where it stores the packets that wait for transmission onto the physical wire.
- When the hardware queue or transmit ring fills, the interface provides explicit back pressure to the L3 processor system. The interface notifies the L3 processor to stop dequeuing packets to the interface transmit ring because the transmit ring is full. The L3 processor now stores the excess packets in the L3 queues.
- When the interface sends the packets on the transmit ring and empties the ring, it once again has sufficient buffers available to store the packets. It releases the back pressure, and the L3 processor dequeues new packets to the interface.

The most important aspect of this communication system is that the interface recognizes that its transmit ring is full and throttles the receipt of new packets from the L3 processor system. Thus, when the interface is congested, the drop decision is moved from a random, last-in/first-dropped decision in the transmit ring first in, first out (FIFO) queue to a differentiated decision based on IP-level service policies implemented by the L3 processor.

What Is the Difference Between "Packets" and "Packets Matched"?

Next, you need to understand when your router uses the L3 queues, since service policies apply only to packets stored in the layer-3 queues.

This table illustrates when packets sit in the L3 queue. Locally generated packets are always process-switched and are delivered first to the L3 queue before they are passed on to the interface driver. Fast-switched and Cisco Express Forwarding (CEF)-switched packets are delivered directly to the transmit ring and sit in the L3 queue only when the transmit ring is full.

Packet Type	Congestion	Non-Congestion
Locally-generated packets, which includes Telnet packets and pings	Yes	Yes
Other packets that are process-switched	Yes	Yes
Packets that are CEF- or fast-switched	Yes	No

This example shows these guidelines applied to the **show policy-map interface** output (the four key counters are in bold):

```
7206# show policy-map interface atm 1/0.1
ATM1/0.1: VC 0/100 -
Service-policy output: cbwfq (1283)
  Class-map: A (match-all) (1285/2)
    28621 packets, 7098008 bytes
    5 minute offered rate 10000 bps, drop rate 0 bps
    Match: access-group 101 (1289)
    Weighted Fair Queueing
      Output Queue: Conversation 73
      Bandwidth 500 (kbps) Max Threshold 64 (packets)
      (pkts matched/bytes matched) 28621/7098008
      (depth/total drops/no-buffer drops) 0/0/0
  Class-map: B (match-all) (1301/4)
    2058 packets, 148176 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: access-group 103 (1305)
    Weighted Fair Queueing
      Output Queue: Conversation 75
      Bandwidth 50 (kbps) Max Threshold 64 (packets)
      (pkts matched/bytes matched) 0/0
      (depth/total drops/no-buffer drops) 0/0/0
  Class-map: class-default (match-any) (1309/0)
    19 packets, 968 bytes
    5 minute offered rate 0 bps, drop rate 0 bps
    Match: any (1313)
```

This table defines the bolded counters:

Counter	Explanation
28621 packets, 7098008 bytes	The number of packets which match the criteria of the class. This counter increments whether or not the interface is congested.
(pkts matched/bytes matched)	The number of packets which match the criteria of the class when the interface was congested. In other words, the interface

28621/7098008	transmit ring was full, and the driver and the L3 processor system worked together to queue the excess packets in the L3 queues, where the service policy applies. Packets that are process-switched always go through the L3 queuing system and thus increment the "packets matched" counter.
Class-map: B (match-all) (1301/4)	These numbers define an internal ID used with the CISCO-CLASS-BASED-QOS-MIB Management Information Base (MIB). They no longer appear in the show policy-map output in current releases of Cisco IOS.
5 minute offered rate 0 bps, drop rate 0 bps	Use the load-interval command to change this value and make it a more instantaneous value. The lowest value is 30 seconds; however, statistics displayed in the show policy-map interface output are updated every 10 seconds. Since the command effectively provides a snapshot at a specific moment, the statistics may not reflect a temporary increase in queue size.

Without congestion, there is no need to queue any excess packets. With congestion, packets, which includes CEF- and fast-switched packets, may go into the L3 queue. Refer back to how the Cisco IOS configuration guide defines congestion: "If you use congestion management features, packets accumulating at an interface are queued until the interface is free to send them; they are then scheduled according to their assigned priority and the queuing mechanism configured for the interface."

Normally, the "packets" counter is much larger than the "pkts matched" counter. If the values of the two counters are nearly equal, then the interface currently receives a large number of process-switched packets or is heavily congested. Both of these conditions should be investigated to ensure optimal packet forwarding.

How Are Conversation Numbers Allocated?

This section explains how your router allocates conversation numbers for the queues created when the service policy is applied.

```
Router# show policy-map interface s1/0.1 dlci 100
Serial1/0.1: DLCI 100 -
output : mypolicy
Class voice
  Weighted Fair Queueing
  Strict Priority
  Output Queue: Conversation 72
    Bandwidth 16 (kbps) Packets Matched 0
    (pkts discards/bytes discards) 0/0
Class immediate-data
  Weighted Fair Queueing
  Output Queue: Conversation 73
    Bandwidth 60 (%) Packets Matched 0
    (pkts discards/bytes discards/tail drops) 0/0/0
  mean queue depth: 0
  drops: class random tail min-th max-th mark-prob
         0      0      0      64    128    1/10
```

```

      1      0      0      71      128      1/10
      2      0      0      78      128      1/10
      3      0      0      85      128      1/10
      4      0      0      92      128      1/10
      5      0      0      99      128      1/10
      6      0      0     106      128      1/10
      7      0      0     113      128      1/10
      rsvp   0      0     120      128      1/10
Class priority-data
  Weighted Fair Queueing
    Output Queue: Conversation 74
      Bandwidth 40 (%) Packets Matched 0 Max Threshold 64 (packets)
      (pkts discards/bytes discards/tail drops) 0/0/0
Class class-default
  Weighted Fair Queueing
    Flow Based Fair Queueing
      Maximum Number of Hashed Queues 64 Max Threshold 20 (packets)

```

The class-default class is the default class to which traffic is directed, if that traffic does not satisfy the match criteria of other classes whose policy is defined in the policy map. The **fair-queue** command allows you to specify the number of dynamic queues into which your IP flows are sorted and classified. Alternately, your router allocates a default number of queues derived from the bandwidth on the interface or VC. Supported values in either case are a power of two, in a range from 16 to 4096.

This table lists the default values for interfaces and for ATM permanent virtual circuits (PVCs):

Bandwidth Range
Number of Dynamic Queues
Less than or equal to 64 kbps
16
More than 64 kbps and less than or equal to 128 kbps
32
More than 128 kbps and less than or equal to 256 kbps
64
More than 256 kbps and less than or equal to 512 kbps
128
More than 512 kbps
256

Bandwidth Range
Number of Dynamic Queues
Less than or equal to 128 kbps

16

More than 128 kbps and less than or equal to 512 kbps

32

More than 512 kbps and less than or equal to 2000 kbps

64

More than 2000 kbps and less than or equal to 8000 kbps

128

More than 8000 kbps

256

Based on the number of reserved queues for weighted fair queuing, Cisco IOS assigns a conversation or queue number as shown in this table:

Conversation / Queue Number	Type of Traffic
1 – 256	General flow-based traffic queues. Traffic that does not match to a user-created class will match to class-default and one of the flow-based queues.
257 – 263	Reserved for Cisco Discovery Protocol (CDP) and for packets marked with an internal high-priority flag.
264	Reserved queue for the priority class (classes configured with the priority command). Look for the "Strict Priority" value for the class in the show policy-map interface output. The priority queue uses a conversation ID equal to the number of dynamic queues plus eight.
265 and higher	Queues for user-created classes.

Confirming Your Service Policy

Complete these steps if you need to test the "pkts matched" counter and your service policy:

1. Simulate congestion with an extended ping using a large ping size and a large number of pings. Also, try to download a large file from a File Transfer Protocol (FTP) server. The file constitutes "disturbing" data and fills the interface bandwidth.
2. Reduce the size of the interface transmit ring with the **tx-ring-limit** command. A reduction of this value accelerates the use of the QoS in the Cisco IOS software.

```
interface ATMx/y.z point-to-point
ip address a.b.c.d M.M.M.M
PVC A/B
tx-ring-limit <size>
```

service-policy output test

3. Specify size as the number of packets for 2600 and 3600 series routers, or as the number of memory particles for 7200 and 7500 series routers.
4. Ensure that your traffic flow matches the input or output parameter of your policy. For example, to download a file from an FTP server generates congestion in the receive direction because the server sends large MTU-sized frames, and the client PC returns small acknowledgments (ACKs).

Related Information

- [QoS Support Page](#)
- [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: Feb 15, 2008

Document ID: 10107
