Understanding the Packet Counters in the show interface rate Command Output with Committed Access Rate (CAR)

Contents

Introduction
Prerequisites
Requirements
Components Used
Conventions
Understanding the show interface rate Command Output
Known Issues with CAR and Class-Based Policing Counters
Related Information

Introduction

Committed Access Rate (CAR) is a rate limiting feature that can be used to provide Classification and Policing services. CAR can be used to classify packets based on certain criteria, such as IP address and port values that use access-lists. The action for packets that conform to the rate limit value and exceed the value can be defined. Refer to Configuring Committed Access Rate for more information on how to configure CAR.

This document explains why the output of the **show interface x/x rate-limit** command shows a `non-zero exceeded bps` value when the `conformed bps` value is less than the configured committed information rate (CIR).

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Understanding the show interface rate Command Output

There are three conditions in which you can see non-zero exceeded rates in the output of this command:

- Burst values are set too low to allow a sufficient throughput rate. For example, see Cisco bug ID CSCdw42923 (registered customers only).
- Resolved issue with double accounting in Cisco IOS® software
- Software bug in Cisco IOS

Look at the example output from a virtual-access interface. In this configuration, RADIUS is used in order to assign a rate limit to the dynamically created virtual-access interface.

```
AV Pair from Radius
Cisco-AVPair = "lcp:interface-config#1=rate-limit input 256000 7500 7500
conform-action continue
exceed-action drop",
Cisco-AVPair = "lcp:interface-config#2=rate-limit output 512000 7500 7500
conform-action continue
exceed-action drop",
```

Use the **show interface x rate-limit** command in order to monitor the performance of the Cisco legacy policer, CAR. In this example, the output of this command provides hints as to why there is a non-zero exceeded bps. The current burst value is 7392 bytes, while the committed burst (Bc) value, indicated by the limit value, is set to 7500 bytes.

```
router#show interfaces virtual-access 26 rate-limit
Virtual-Access26 Cable Customers
  Input
    matches: all traffic
      params:  256000 bps, 7500 limit, 7500 extended limit
      conformed 2248 packets, 257557 bytes; action: continue
      exceeded 35 packets, 22392 bytes; action: drop
      last packet: 156ms ago, current burst: 0 bytes
      last cleared 00:02:49 ago, conformed 12000 bps, exceeded 1000 bps
  Output
    matches: all traffic
      params:  512000 bps, 7500 limit, 7500 extended limit
      conformed 3338 packets, 4115194 bytes; action: continue
      exceeded 565 packets, 797648 bytes; action: drop
      last packet: 188ms ago, current burst: 7392 bytes
      last cleared 00:02:49 ago, conformed 194000 bps, exceeded 37000 bps
```

When you configure CAR or a newer policer from Cisco, class-based policing, you must configure sufficiently high burst values in order to assure expected throughput and in order to ensure that the policer drops packets only to punish short-term congestion.

When you select burst values, it is important to accommodate transient increases in the queue size. You cannot simply assume that packets arrive and depart at the same time. You also cannot assume that the queue changes from empty to one packet and that the queue remains at one packet based on a consistent one in/one out arrival time. If the typical traffic is fairly bursty, then the burst values need to be correspondingly large in order to allow the link utilization to be maintained at an acceptably high level. A burst size that is too low, or a minimum threshold that is

too low, can result in unacceptably low link utilization.

A burst can be defined simply as a series of back-to-back, MTU-sized frames, such as 1500-byte frames that originate on an Ethernet network. When a burst of such frames arrives at an output interface, it can overwhelm the output buffers and exceed the configured depth of the token bucket at an instantaneous moment in time. With the use of a token metering system, a policer makes a binary decision about whether an arriving packet conforms, exceeds, or violates the configured policing values. With bursty traffic, such as an FTP stream, the instantaneous arrival rate of these packets can exceed the configured burst values and lead to CAR drops.

In addition, overall throughput in times of congestion vary with the type of traffic that is evaluated by the policer. While TCP traffic is responsive to congestion, other flows are not. Examples of non-responsive flows include UDP-based and ICMP-based packets.

TCP is based on positive acknowledgement with retransmission. TCP uses a sliding window as part of its positive acknowledgement mechanism. Sliding window protocols use network bandwidth better because they allow the sender to transmit multiple packets before they wait for an acknowledgement. For example, in a sliding window protocol with a window size of 8, the sender is permitted to transmit 8 packets before it receives an acknowledgement. If you increase the window size, the network idle time is largely eliminated. A well-tuned sliding window protocol keeps the network completely saturated with packets and maintains high throughput.

Since endpoints do not know the specific congestion status of the network, TCP as a protocol is designed react to congestion in the network by the reduction its transmission rates when congestion occurs. Specifically, it uses two techniques:

Technique	**Description**
Multiplicative decrease congestion avoidance	Upon loss of a segment (the equivalent of a packet to TCP), reduce the congestion window by half. The congestion window is a second value or window which is used to limit the number of packets that a sender can transmit into the network before it waits for an acknowledgement.
Slow start recovery	When you start traffic on a new connection or increase traffic after a period of congestion, start the congestion window at the size of a single segment and increase the congestion window by one segment each time an acknowledgement arrives. TCP initializes the congestion window to 1, sends an initial segment, and waits. When the acknowledgement arrives, it increases the congestion window to 2, sends two segments, and waits. For more details, see RFC 2001 .

Packets can be lost or destroyed when transmission errors interfere with data, when network hardware fails, or when networks become too heavily loaded to accommodate the load presented. TCP assumes that lost packets, or packets that fail to be acknowledged within the timed interval due to extreme delay, indicate congestion in the network.

The token-bucket metering system of a policer is invoked on each packet arrival. Specifically, the conformed rate and exceed rate are calculated based on this simple formula:

```
router#show interfaces virtual-access 26 rate-limit
Virtual-Access26 Cable Customers
  Input
    matches: all traffic
      params:  256000 bps, 7500 limit, 7500 extended limit
      conformed 2248 packets, 257557 bytes; action: continue
      exceeded 35 packets, 22392 bytes; action: drop
      last packet: 156ms ago, current burst: 0 bytes
      last cleared 00:02:49 ago, conformed 12000 bps, exceeded 1000 bps
  Output
    matches: all traffic
      params:  512000 bps, 7500 limit, 7500 extended limit
      conformed 3338 packets, 4115194 bytes; action: continue
      exceeded 565 packets, 797648 bytes; action: drop
      last packet: 188ms ago, current burst: 7392 bytes
      last cleared 00:02:49 ago, conformed 194000 bps, exceeded 37000 bps
```

Since the formula calculates rates over a period from the last time that the counters were cleared, Cisco recommends to clear the counters in order to monitor the current rate. If the counters are not cleared, then the previous formula rate effectively means that the **show** command output displays an average calculated over a potentially very long period, and the values possibly are not meaningful in the determination of the current rate.

The average throughput should match the configured committed information rate (CIR) over a period of time. Burst sizes allow a maximum burst duration at a given time. If there is no traffic or less than the CIR's worth of traffic and the token bucket does not fill, a very large burst is still limited to a particular size calculated based on normal burst and extended burst.

The drop rate results from this mechanism

1. Note the current time.
2. Update the token bucket with the number of tokens that have accumulated continuously since the last time a packet arrived.
3. The total number of accumulated tokens cannot exceed the maxtokens value. Drop excess tokens.
4. Check for packet conformance.

Rate-limiting can also be achieved with Policing. This is a sample configuration in order to provide rate-limiting on the Ethernet interface that uses Class based policing.

```
router#show interfaces virtual-access 26 rate-limit
Virtual-Access26 Cable Customers
  Input
    matches: all traffic
      params:  256000 bps, 7500 limit, 7500 extended limit
      conformed 2248 packets, 257557 bytes; action: continue
      exceeded 35 packets, 22392 bytes; action: drop
      last packet: 156ms ago, current burst: 0 bytes
      last cleared 00:02:49 ago, conformed 12000 bps, exceeded 1000 bps
  Output
    matches: all traffic
      params:  512000 bps, 7500 limit, 7500 extended limit
```

```
        conformed 3338 packets, 4115194 bytes; action: continue
        exceeded 565 packets, 797648 bytes; action: drop
        last packet: 188ms ago, current burst: 7392 bytes
        last cleared 00:02:49 ago, conformed 194000 bps, exceeded 37000 bps
```

This sample output from the **show policy-map interface** command illustrates properly calculated and synchronized values for offered rate and drop rate as well as conformed and exceed bps rates.

```
router#show policy-map interface ethernet 3/0
  Ethernet3/0

   Service-policy input: p2

     Class-map: rtp1 (match-all)
       88325 packets, 11040625 bytes
       30 second offered rate 400000 bps, drop rate 150000 bps
       Match: ip rtp 2000 10
       police:
         250000 bps, 7750 limit, 7750 extended limit
         conformed 55204 packets, 6900500 bytes; action: transmit
         exceeded 33122 packets, 4140250 bytes; action: drop
         conformed 250000 bps, exceed 150000 bps violate 0 bps

       Service-policy : p3b

         Class-map: rtp1 (match-all)
           88325 packets, 11040625 bytes
           30 second offered rate 400000 bps, drop rate 50000 bps
           Match: ip rtp 2000 10
           police:
             200000 bps, 6250 limit, 6250 extended limit
             conformed 44163 packets, 5520375 bytes; action: transmit
             exceeded 11041 packets, 1380125 bytes; action: drop
             conformed 200000 bps, exceed 50000 bps violate 0 bps

         Class-map: class-default (match-any)
           0 packets, 0 bytes
           30 second offered rate 0 bps, drop rate 0 bps
           Match: any
```

Known Issues with CAR and Class-Based Policing Counters

This table lists resolved issues with the counters displayed in the **show policy-map** or **show interface rate-limit** commands. Registered customers who are logged in can view the bug information in the Bug Search Tool.

Symptom	Resolved Bug IDs and Workarounds
Lower than expected drop counters	• Cisco bug ID CSCdv41231 (registered customers only)When an input hierarchical service policy uses the **police** command at the parent and child levels, the policer can drop less than the expected number of packets since the parent-level policer must be

	congested before it drops the packets. This is an example of such a policy:
	`router#show policy-map interface ethernet 3/0`
	`  Ethernet3/0`
	
	`   Service-policy input: p2`
	
	`     Class-map: rtp1 (match-all)`
	`       88325 packets, 11040625 bytes`
	`       30 second offered rate 400000 bps, drop rate 150000 bps`
	`       Match: ip rtp 2000 10`
	`       police:`
	`         250000 bps, 7750 limit, 7750 extended limit`
	`         conformed 55204 packets, 6900500 bytes; action: transmit`
	`         exceeded 33122 packets, 4140250 bytes; action: drop`
	`         conformed 250000 bps, exceed 150000 bps violate 0 bps`
	
	`       Service-policy : p3b`
	
	`         Class-map: rtp1 (match-all)`
	`           88325 packets, 11040625 bytes`
	`           30 second offered rate 400000 bps, drop rate 50000 bps`
	`           Match: ip rtp 2000 10`
	`           police:`
	`             200000 bps, 6250 limit, 6250 extended limit`
	`             conformed 44163 packets, 5520375 bytes; action: transmit`
	`             exceeded 11041 packets, 1380125 bytes; action: drop`
	`             conformed 200000 bps, exceed 50000 bps violate 0 bps`
	
	`         Class-map: class-default (match-any)`
	`           0 packets, 0 bytes`
	`           30 second offered rate 0 bps, drop rate 0 bps`
	`           Match: any`
	As a workaround, create separate policies and apply one on inbound and one on outbound in order to avoid the configuration of a hierarchical policy.

Double the expected rate of drops and throughput	• Cisco bug ID CSCds23924 (registered customers only)Cisco Express Forwarding (CEF) defines an IOS switching mechanism which forwards packets from input to output interface. Prior to the changes implemented from this bug ID, both CEF and configured QoS mechanisms such as CAR or class-based policing incremented the packet counters. The result is so-called double accounting and inflated conformed packets

<table>
<tr><td></td><td>and excess drop values.
• Cisco bug ID CSCdr40598 (registered customers only)On the Cisco 12000 series, when output CAR is enabled and the ingress line card is Engine 2, the egress output counters are doubled. This double accounting results from how output counters are handled.
• Cisco bug ID CSCdv84259 (registered customers only)If you globally enable the ip cef distributed command on a Cisco 7500 series router, a non-Versatile Interface Processor (VIP) card interface appears with the ip route-cache distributed command enabled by default. Non-VIPs do not support distributed CEF, and a rare side-effect of this command that appears on non-VIPs is double accounting.</td></tr>
<tr><td>No drops or a zero drop rate</td><td>In general, when you apply class-based QoS features, the first step in troubleshooting is to ensure that the QoS classification mechanism works properly. In other words, ensure that the packets specified in the match statements in your class-map hit the correct classes.
<pre>router#show policy-map interface
 ATM4/0.1

  Service-policy input: drop-inbound-http-hacks
(1061)

    Class-map: http-hacks (match-any) (1063/2)
      149 packets, 18663 bytes
      5 minute offered rate 2000 bps, drop rate 0
bps
      Match: protocol http url "*cmd.exe*" (1067)
        145 packets, 18313 bytes
        5 minute rate 2000 bps
      Match: protocol http url "*.ida*" (1071)
        0 packets, 0 bytes
        5 minute rate 0 bps
      Match: protocol http url "*root.exe*"
(1075)
        4 packets, 350 bytes
        5 minute rate 0 bps
      Match: protocol http url "*readme.eml*"
(1079)
        0 packets, 0 bytes
        5 minute rate 0 bps
      police:
        1000000 bps, 31250 limit, 31250 extended
limit
        conformed 0 packets, 0 bytes; action:
drop
        exceeded 0 packets, 0 bytes; action: drop
        violated 0 packets, 0 bytes; action: drop
        conformed 0 bps, exceed 0 bps violate 0</pre></td></tr>
</table>

<table>
<tr><td></td><td>bps
Cisco bug ID CSCds34478 (registered customers only)Classification fails when CEF, and not DCEF, is enabled and an input policy is attached to an ATM PVC. In Cisco IOS Software Release 12.1T, output classification fails when CEF, and not DCEF, is enabled and an output policy is attached to an ATM PVC.</td></tr>
<tr><td>Anomalous or inconsistent drop rate</td><td>Cisco bug ID CSCdw50583 (registered customers only)The drop rate displayed in the class-map does not match the drop rates indicated by the police action. In this example output, the drop rate for the class is 745000 bps, while the drop rate shown by police action is 1072000 bps.
router#show policy-map interface
Serial3/0.1: DLCI 13 -

Service-policy output: out

Class-map: c2 (match-all)
172483 packets, 91760956 bytes
30 second offered rate 1384000 bps, drop rate 745000 bps
Match: ip precedence 0
police:
384000 bps, 1500 limit, 1500 extended limit
conformed 38903 packets, 20696396 bytes; action: transmit
exceeded 133580 packets, 71064560 bytes; action: drop
conformed 311000 bps, exceed 1072000 bps violate 0 bps</td></tr>
</table>

Related Information

- **Configuring Committed Access Rate**
- **Policing with CAR**
- **Using CAR During DOS Attacks**
- **QoS Technology Support Page**
- **IP Routed Protocols Support Page**
- **IP Routing Support Page**
- **Technical Support & Documentation - Cisco Systems**