Configuring Denial of Service Protection

This chapter contains information on how to protect your switch against Denial of Service (DoS) attacks. The information covered in this chapter is unique to Cisco IOS Release 12.2SX, and it supplements the network security information and procedures in the "Configuring Network Security" chapter in this publication as well as the network security information and procedures in these publications:

- *Cisco IOS Security Configuration Guide*, Release 12.2, at this URL:

  http://www.cisco.com/en/US/docs/ios/12_2/security/configuration/guide/fsecur_c.html

- *Cisco IOS Security Command Reference*, Release 12.2, at this URL:

  http://www.cisco.com/en/US/docs/ios/12_2/security/command/reference/fsecur_r.html

Note For complete syntax and usage information for the commands used in this chapter, see these publications:

- The Cisco IOS Master Command List, at this URL:

  http://www.cisco.com/en/US/docs/ios/mcl/allreleasemcl/all_book.html

- The Release 12.2 publications at this URL:

  http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_installation_and_configuration_guides_list.html

Tip For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

Participate in the Technical Documentation Ideas forum

This chapter consists of these sections:

- Understanding DoS Protection, page 52-2
- DoS Protection Default Configuration, page 52-13
- DoS Protection Configuration Guidelines and Restrictions, page 52-13
- Configuring Sticky ARP, page 52-18

Understanding DoS Protection

This section contains information about the available methods to counteract DoS attacks and includes configuration examples. The PFC3 provides a layered defense against DoS attacks using the following methods:

- CPU rate limiters—Controls traffic types.
- Control plane policing (CoPP)—Filters and rate limits control plane traffic. For information about CoPP, see Chapter 53, "Configuring Control Plane Policing."

These sections describe DoS protection:

Security ACLs and VACLs

If the network is under a DoS attack, ACLs can be an efficient method for dropping the DoS packets before they reach the intended target. Use security ACLs if an attack is detected from a particular host.

In this example, the host 10.1.1.10 and all traffic from that host is denied:

```
Router(config)# access-list 101 deny ip host 10.1.1.10 any
Router(config)# access-list 101 permit ip any any
```

Security ACLs also protect against the spoofing of addresses. For example, assume that a source address A is on the inside of a network and a switch interface that is pointing to the Internet. You can apply an inbound ACL on the switch Internet interface that denies all addresses with a source of A (the inside address). This action stops attacks where the attackers spoof inside source addresses. When the packet arrives at the switch interface, it matches on that ACL and drops the packet before it causes damage.

When the switch is used with a Cisco Intrusion Detection Module (CIDM), you can dynamically install the security ACL as a response to the detection of the attack by the sensing engine.

VACLs are a security enforcement tool based on Layer 2, Layer 3, and Layer 4 information. The result of a VACL lookup against a packet can be a permit, a deny, a permit and capture, or a redirect. When you associate a VACL with a particular VLAN, all traffic must be permitted by the VACL before the traffic is allowed into the VLAN. VACLs are enforced in hardware, so there is no performance penalty for applying VACLs to a VLAN.

QoS Rate Limiting

QoS ACLs limit the amount of a particular type of traffic that is processed by the RP. If a DoS attack is initiated against the RP, QoS ACLs can prevent the DoS traffic from reaching the RP data path and congesting it. The PFC3 performs QoS in hardware, which offers an efficient means of limiting DoS traffic (once that traffic has been identified) to protect the switch from impacting the RP.

For example, if the network is experiencing ping-of-death or smurf attacks, the administrator should rate limit the ICMP traffic to counteract the DoS attack and still allow legitimate traffic through the processor, or allow it to be forwarded to the RP or host. This rate limiting configuration must be done for each flow that should be rate limited and the rate-limiting policy action should be applied to the interface.

In the following example, the access-list 101 permits and identifies ping (echo) ICMP messages from any source to any destination as traffic. Within the policy map, a policing rule defines a specified committed information rate (CIR) and burst value (96000 bps and 16000 bps) to rate limit the ping (ICMP) traffic through the chassis. The policy map then is applied to an interface or VLAN. If the ping traffic exceeds the specified rate on the VLAN or interface where the policy map is applied, it is dropped as specified in the markdown map (the markdown map for the normal burst configurations is not shown in the example).

```
Router(config)# access-list 101 permit icmp any any echo
Router(config)# class-map match-any icmp_class
Router(config-cmap)# match access-group 101
Router(config-cmap)# exit
Router(config)# policy-map icmp_policer
Router(config-pmap)# class icmp_class
Router(config-pmap-c)# police 96000 16000 conform-action transmit exceed-action
policed-dscp-transmit drop
Router(config-pmap-c)# exit
Router(config-pmap)# exit
```

uRPF Check

When you enable the unicast reverse path forwarding (uRPF) check, packets that lack a verifiable source IP address, such as spoofed IP source addresses, are discarded. Cisco Express Forwarding (CEF) tables are used to verify that the source addresses and the interfaces on which they were received are consistent with the switch processor (SP) FIB tables.

After you enable uRPF check on an interface (per-VLAN basis), the incoming packet is compared to the CEF tables through a reverse lookup. If the packet is received from one of the reverse path routes, the packet is forwarded. If there is no reverse path route on the interface on which the packet was received, the packet fails the uRPF check and is either dropped or forwarded, depending on whether an ACL is applied to the uRPF check fail traffic. If no ACL is specified in the CEF tables, then the forged packets are immediately dropped.

You can only specify an ACL for the uRPF check for packets that fail the uRPF check. The ACL checks whether the packet should immediately be dropped or forwarded. The uRPF check with ACL is not supported in any PFC3 in hardware. Packets that are denied in the uRPF ACL are forwarded in hardware. Packets that are permitted are sent to the CPU.

The uRPF check is supported in hardware. However, all packets that fail the uRPF check, and are forwarded because of an applied ACL, can be sent and rate limited to the RP to generate ICMP unreachable messages; these actions are all software driven. The uRPF check in hardware is supported for routes with up to two return paths (interfaces) and up to six return paths with interface groups configured (two from the FIB table and four from the interface groups).

Traffic Storm Control

A traffic storm occurs when packets flood the LAN, which creates excessive traffic and degrades network performance. The traffic storm control feature prevents LAN ports from being disrupted by a broadcast, multicast, or unicast traffic storm on physical interfaces from either mistakes in network configurations or from users issuing a DoS attack. Traffic storm control (also called traffic suppression)

monitors incoming traffic levels over a 1-second traffic storm control interval. During the interval, traffic storm control compares the traffic level with the configured traffic storm control level. The traffic storm control level is a percentage of the total available bandwidth of the port. Each port has a single traffic storm control level that is used for all types of traffic (broadcast, multicast, and unicast).

Traffic storm control is configured on an interface and is disabled by default. The configuration example here enables broadcast address storm control on interface FastEthernet 2/3 to a level of 20 percent. When the broadcast traffic exceeds the configured level of 20 percent of the total available bandwidth of the port within a 1-second traffic-storm-control interval, traffic storm control will drop all broadcast traffic until the end of the traffic-storm-control interval.

```
Router(config-if)# storm-control broadcast level 20
```

The switch supports broadcast storm control on all LAN ports and multicast and unicast storm control on Gigabit Ethernet ports.

When two or three suppression modes are configured simultaneously, they share the same level settings. If broadcast suppression is enabled, and if multicast suppression is also enabled and configured at a 70-percent threshold, the broadcast suppression will also have a setting for 70 percent.

Network Under SYN Attack

A network under a SYN attack is easily recognized. The target host becomes unusually slow, crashes, or suspends operation. Traffic returned from the target host can also cause trouble on the RP because return traffic goes to randomized source addresses of the original packets, lacks the locality of "real" IP traffic, and may overflow route caches, or CEF tables.

When the network is under a SYN attack, the TCP intercept feature becomes aggressively defensive. Two factors determine when aggressive behavior on the switch begins and ends:

*   The total incomplete connections
*   Connection requests during the last one-minute sample period

Both factors are configured with low and high values.

If the number of incomplete connections exceed 1,100, or the number of connections arriving in the last one-minute period exceed 1,100, each new arriving connection causes the oldest partial connection (or a random connection) to be deleted. These are the default values, which can be altered. When either of the thresholds is exceeded, the TCP intercept assumes the server is under attack and goes into aggressive mode with the following reactions:

*   Each new arriving connection causes the oldest partial (or random partial) to be deleted.

*   The initial retransmission timeout is reduced by half to 0.5 seconds, and so the total time trying to establish the connection is cut in half.

*   In watch mode, the watch timeout is reduced by half.

Note   When both thresholds fall below the configured low value, the aggressive behavior ceases (default value is 900 in both factors).

TCP flows are hardware assisted on the PFC3.

Protocol Packet Policing

Attackers may try to overwhelm the RP CPU with routing protocol or ARP control packets. The **mls qos protocol** command rate limits this traffic in hardware. Enter the **mls qos protocol ?** to display the supported routing protocols.

For example, the **mls qos protocol arp police** command rate limits ARP packets.

ARP packets are approximately 40 bytes long and ARP reply packets are approximately 60 bytes long. The policer rate value is in bits per second. The burst value is in bytes per second. Together, an ARP request and reply are approximately 800 bits.

This example shows how to allow 200 ARP requests and replies per second:

```
Router(config)# mls qos protocol arp police 200000 6000
```

Note
- The minimum values supported by the **mls qos protocol arp police** command are too small for use in production networks.

- The configured rate limits are applied separately to the PFC and each DFC. The RP CPU will receive the configured value times the number of forwarding engines.

- The protocol packet policing mechanism effectively protects the RP CPU against attacks such as line-rate ARP attacks, but it polices both routing protocols and ARP packets to the switch and also polices traffic through the switch with less granularity than CoPP.

- The policing mechanism shares the root configuration with a policing-avoidance mechanism. The policing-avoidance mechanism lets the routing protocol and ARP packets flow through the network when they reach a QoS policer. This mechanism can be configured using the **mls qos protocol** *protocol_name* **pass-through** command.

This example shows how to display the available protocols to use with protocol packet policing:

```
Router(config)# mls qos protocol ?
  isis
  eigrp
  ldp
  ospf
  rip
  bgp
  ospfv3
  bgpv2
  ripng
  neigh-discover
  wlccp
  arp
```

This example shows how to display the available keywords to use with the **mls qos protocol** command:

```
Router(config)# mls qos protocol protocol_name ?
  pass-through  pass-through keyword
  police        police keyword
  precedence    change ip-precedence(used to map the dscp to cos value)
```

Recommended Rate-Limiter Configuration

The recommended rate-limiter configuration is as follows:

- Enable the rate limiters for the traffic types most likely to be used in a DoS attack.
- Do not use a rate limiter on VACL logging unless you configure VACL logging.
- Disable redirects.
- Disable unreachables.
- Do not enable the MTU rate limiter if all interfaces have the same MTU.
- When configuring the Layer 2 PDU rate limiter, note the following information:
  - Calculate the expected or possible number of valid PDUs and double or triple the number.
  - PDUs include BPDUs, DTP, VTP, PAgP, LACP, UDLD, etc.
  - Rate limiters do not discriminate between good frames or bad frames.

Hardware-Based Rate Limiters on the PFC3

The PFC3 supports additional hardware-based rate limiters. The PFC3 provides eight rate-limiter registers for the new rate limiters, which are configured globally on the switch. These rate-limiter registers are present in the Layer 3 forwarding engine (PFC) and are responsible for containing rate-limiting information for result packets that match the various available configured rate limiters.

Because eight rate-limiter registers are present on the PFC3, these registers can force different rate-limiting scenarios to share the same register. The registers are assigned on a first-come, first-serve basis. If all registers are being utilized, the only way to configure another rate limiter is to free one register.

These are the hardware-based rate limiters available on the PFC3:

- Ingress-Egress ACL Bridged Packets (Unicast Only), page 52-7
- uRPF Check Failure, page 52-7
- TTL Failure, page 52-8
- ICMP Unreachable (Unicast Only), page 52-8
- FIB (CEF) Receive Cases (Unicast Only), page 52-8
- FIB Glean (Unicast Only), page 52-9
- Layer 3 Security Features (Unicast Only), page 52-9
- ICMP Redirect (Unicast Only), page 52-9
- VACL Log (Unicast Only), page 52-9
- MTU Failure, page 52-10
- Layer 2 PDU, page 52-10
- Layer 2 Protocol Tunneling, page 52-10
- IP Errors, page 52-11
- Layer 2 Multicast IGMP Snooping, page 52-10
- IPv4 Multicast, page 52-11
- IPv6 Multicast, page 52-12

Ingress-Egress ACL Bridged Packets (Unicast Only)

This rate limiter rate limits packets sent to the RP because of an ingress/egress ACL bridge result. The switch accomplishes this by altering existing and new ACL TCAM entries with a TCAM bridge result to a Layer 3 redirect result pointing to the RP. Packets hitting the TCAM entries with the altered Layer 3 redirect rate limit result will be rate limited according to the instructions set in CLI by the network administrator. Both the ingress and egress values will be the same, as they both share the same rate-limiter register. If the ACL bridge ingress/egress rate limiting is disabled, the Layer 3 redirect rate limit results are converted to the bridge result.

Ingress or egress ACL-bridged packet cases share a single rate-limiter register. If the feature is turned on, ingress and egress ACLs use the same rate-limiter value.

Burst values regulate how many packets can be allowed in a burst. Each allowed packet consumes a token and a token must be available for a packet to be allowed. One token is generated per millisecond. When packets are not coming in, tokens can be accumulated up to the burst value. For example, if the burst value is set to 50, the switch can accumulate up to 50 tokens and absorb a burst of 50 packets.

This example shows how to rate limit the unicast packets from an ingress ACL bridge result to 50000 packets per second, and 50 packets in burst:

```
Router(config)# mls rate-limit unicast acl input 50000 50
```

This example shows how to rate limit the unicast packets from an ingress ACL bridge result to the same rate (50000 pps and 50 packets in burst) for egress ACL bridge results:

```
Router(config)# mls rate-limit unicast acl output 50000 50
```

If the values of the rate limiter are altered on either the ingress or the egress when both are enabled, both values are changed to that new value. In the following example, the output rate is changed to 40000 pps:

```
Router(config)# mls rate-limit unicast acl output 40000 50
```

When you enter the **show mls rate-limit** command, both the ACL bridged in and the ACL bridged out display the new value of 40000 pps:

```
Router# show mls rate-limit
  Rate Limiter Type        Status      Packets/s   Burst
--------------------    ----------   ---------   -----
MCAST NON RPF           Off          -           -
MCAST DFLT ADJ          On           100000      100
MCAST DIRECT CON        Off          -           -
ACL BRIDGED IN          On           40000       50
ACL BRIDGED OUT         On           40000       50
IP FEATURES             Off
…
```

uRPF Check Failure

The uRPF check failure rate limiter allows you to configure a rate for the packets that need to be sent to the RP because they failed the uRPF check. The uRPF checks validate that incoming packets on an interface are from a valid source, which minimizes the potential threat of DoS attacks from users using spoofed addresses. When spoofed packets fail the uRPF check, those failures can be sent to the RP. The uRPF check rate limiters allow you to rate limit the packets per second that are bridged to the RP CPU when a uRPF check failure occurs.

This example shows how to rate limit the uRPF check failure packets sent to the RP to 100000 pps with a burst of 100 packets:

```
Router(config)# mls rate-limit unicast ip rpf-failure 100000 100
```

TTL Failure

This rate limiter rate limits packets sent to the RP because of a time-to-live (TTL) check failure. As indicated by the **all** keyword in the following example, this rate limiter applies to both multicast and unicast traffic.

✎

Note   The TTL failure rate limiter is not supported for IPv6 multicast.

This example shows how to rate limit the TTL failures to 70000 pps with a burst of 150:

```
Router(config)# mls rate-limit all ttl-failure 70000 150
```

For packets having TTL value equal to 1, configure MLS ttl-failure rate-limiter for CPU protection. The following are a few exceptions:

- When the incoming traffic also requires ARP resolution (known network, unknown host), use glean rate-limiter.

- When the incoming traffic is also destined to the IP of the router, use receive-rate limiter.

- When the incoming traffic is also receiving input ACL bridged (for example, ACL log), use acl bridge input rate-limiter.

ICMP Unreachable (Unicast Only)

In an ICMP unreachable attack, a device is flooded with a large number of packets that contain a destination address that is unreachable from the flooded device (in this case, the RP). The ICMP unreachable rate limiter allows you to rate limit the packets that are sent to the RP containing unreachable addresses.

This example shows how to rate limit the packets that are sent to the RP because of an ACL drop to 10000 pps and a burst of 100:

```
Router(config)# mls rate-limit unicast ip icmp unreachable acl-drop 10000 100
```

This example shows how to rate limit the packets that require generation of ICMP-unreachable messages because of a FIB miss to 80000 pps and burst to 70:

```
Router(config)# mls rate-limit unicast ip icmp unreachable no-route 80000 70
```

The four rate limiters, ICMP unreachable no route, ICMP unreachable ACL drop, IP errors, and IP RPF failure, share a single rate-limiter register. If any of these limiters are enabled, all of the limiters in this group will share the same value and sometimes the same state (for example, ON/ON/ON). When verifying the rate limiters, if the members of this register are enabled through another feature, an ON-Sharing status (instead of an ON status) is displayed. The exception is the TTL failure rate limiter: its value shares the same value as the other members in the register if you have manually enabled the feature.

FIB (CEF) Receive Cases (Unicast Only)

The FIB receive rate limiter provides the capability to rate limit all packets that contain the RP IP address as the destination address. The rate limiters do not discriminate between good frames and bad frames.

✎

Note   Do not enable the FIB receive rate limiter if you are using CoPP. The FIB receive rate limiter overrides the CoPP policies.

This example shows how to rate limit the traffic to 25000 pps with a burst of 60:

```
Router(config)# mls rate-limit unicast cef receive 25000 60
```

FIB Glean (Unicast Only)

The FIB glean rate limiter does not limit ARP traffic, but provides the capability to rate limit traffic that requires address resolution (ARP) and requires that it be sent to the RP. This situation occurs when traffic enters a port and contains the destination of a host on a subnet that is locally connected to the RP, but no ARP entry exists for that destination host. In this case, because the MAC address of the destination host will not be answered by any host on the directly connected subnet that is unknown, the "glean" adjacency is hit and the traffic is sent directly to the RP for ARP resolution. This rate limiter limits the possibility of an attacker overloading the CPU with such ARP requests.

This example shows how to rate limit the rate at which this traffic is sent to the RP to 20000 pps and a burst of 60:

```
Router(config)# mls rate-limit unicast cef glean 20000 60
```

Layer 3 Security Features (Unicast Only)

Some security features are processed by first being sent to the RP. For these security features, you need to rate limit the number of these packets being sent to the RP to reduce any potential overloading. The security features include authentication proxy (auth-proxy), IPSEC, and inspection.

Authentication proxy is used to authenticate inbound or outbound users or both. These users are normally blocked by an access list, but with auth-proxy, the users can bring up a browser to go through the firewall and authenticate on a terminal access controller access control system plus (TACACS+) or RADIUS server (based on the IP address). The server passes additional access list entries down to the switch to allow the users through after authentication. These ACLs are stored and processed in software, and if there are many users utilizing auth-proxy, the RP may be overwhelmed. Rate limiting would be advantageous in this situation.

IPsec and inspection are also done by the RP and may require rate limiting. When the Layer 3 security feature rate limiter is enabled, all Layer 3 rate limiters for auth-proxy, IPsec and inspection are enabled at the same rate.

This example shows how to rate limit the security features to the RP to 100000 pps with a burst of 10 packets:

```
Router(config)# mls rate-limit unicast ip features 100000 10
```

ICMP Redirect (Unicast Only)

The ICMP-redirect rate limiter allows you to rate limit ICMP traffic. For example, when a host sends packets through a nonoptimal switch, the RP sends ICMP-redirect messages to the host to correct its sending path. If this traffic occurs continuously, and is not rate limited, the RP will continuously generate ICMP-redirect messages.

This example shows how to rate limit the ICMP redirects to 20000 pps, with a burst of 20 packets:

```
Router(config)# mls rate-limit unicast ip icmp redirect 20000 20
```

VACL Log (Unicast Only)

Packets that are sent to the RP because of VLAN-ACL logging can be rate limited to ensure that the CPU is not overwhelmed with logging tasks. VACLs are processed in hardware, but the RP does the logging. When VACL logging is configured on the switch, IP packets that are denied in the VACL generate log messages.

This example shows how to rate limit logging requests to 5000 pps (the range for this rate limiter is from 10 to 5000 pps):

```
Router(config)# mls rate-limit unicast acl vacl-log 5000
```

MTU Failure

Similar to the TTL failure rate limiter, the rate limiter for MTU failures is supported for both unicast and multicast traffic. Packets that fail an MTU check are sent to the RP CPU. This might cause the RP to be overwhelmed.

This example shows how to rate limit packets failing the MTU failures from being sent to the RP to 10000 pps with a burst of 10:

```
Router(config)# mls rate-limit all mtu 10000 10
```

Layer 2 Multicast IGMP Snooping

The IGMP snooping rate limiter limits the number of Layer 2 IGMP packets destined for the SP. IGMP snooping listens to IGMP messages between the hosts and the switch. You cannot enable the Layer 2 PDU rate limiter if the switch is operating in truncated mode. The switch uses truncated mode for traffic between fabric-enabled modules when there are both fabric-enabled and nonfabric-enabled modules installed. In this mode, the switch sends a truncated version of the traffic (the first 64 bytes of the frame) over the switch fabric channel. In releases earlier than Cisco IOS Release 12.2(33)SXH, the IGMP snooping rate limiter also rate limits PIM messages. In Cisco IOS Release 12.2(33)SXH and later releases, the IGMP snooping rate limiter does not rate limit PIM messages.

This example shows how to rate limit IGMP-snooping traffic:

```
Router(config)# mls rate-limit multicast ipv4 igmp 20000 40
```

Layer 2 PDU

The Layer 2 protocol data unit (PDU) rate limiter allows you to limit the number of Layer 2 PDU protocol packets (including BPDUs, DTP, PAgP, CDP, STP, and VTP packets) destined for the SP and not the RP CPU. You cannot enable the Layer 2 PDU rate limiter if the switch is operating in truncated mode. The switch uses truncated mode for traffic between fabric-enabled modules when there are both fabric-enabled and nonfabric-enabled modules installed. In this mode, the switch sends a truncated version of the traffic (the first 64 bytes of the frame) over the switch fabric channel.

This example shows how to rate limit Layer 2 PDUs to 20000 pps with a burst of 20 packets:

```
Router(config)# mls rate-limit layer2 pdu 20000 20
```

Layer 2 Protocol Tunneling

This rate limiter limits the Layer 2 protocol tunneling packets, which include control PDUs, CDP, STP, and VTP packets destined for the SP. These packets are encapsulated in software (rewriting the destination MAC address in the PDU), and then forwarded to a proprietary multicast address (01-00-0c-cd-cd-d0). You cannot enable the Layer 2 PDU rate limiter if the switch is operating in truncated mode. The switch uses truncated mode for traffic between fabric-enabled modules when there are both fabric-enabled and nonfabric-enabled modules installed. In this mode, the switch sends a truncated version of the traffic (the first 64 bytes of the frame) over the switch fabric channel.

This example shows how to rate limit Layer 2 protocol tunneling packets to 10000 pps with a burst of 10 packets:

```
Router(config)# mls rate-limit layer2 l2pt 10000 10
```

IP Errors

This rate limiter limits the packets with IP checksum and length errors. When a packet reaches the PFC3 with an IP checksum error or a length inconsistency error, it must be sent to the RP for further processing. An attacker might use the malformed packets to carry out a DoS attack, but the network administrator can configure a rate for these types of packets to protect the control path.

This example shows how to rate limit IP errors sent to the RP to 1000 pps with a burst of 20 packets:

```
Router(config)# mls rate-limit unicast ip errors 1000 20
```

IPv4 Multicast

This rate limiter limits the IPv4 multicast packets. The rate limiters can rate limit the packets that are sent from the data path in the hardware up to the data path in the software. The rate limiters protect the control path in the software from congestion and drop the traffic that exceeds the configured rate. Within the IPv4 multicast rate limiter, there are three rate limiters that you can also configure: the FIB-miss rate limiter, the multicast partially switched flows rate limiter, and the multicast directly connected rate limiter.

The FIB-miss rate limiter allows you to rate limit the multicast traffic that does not match an entry in the mroute table.

The partially switched flow rate limiter allows you to rate limit the flows destined to the RP for forwarding and replication. For a given multicast traffic flow, if at least one outgoing Layer 3 interface is multilayer switched, and at least one outgoing interface is not multilayer switched (no H-bit set for hardware switching), the particular flow is considered partially switched, or partial-SC (partial shortcut). The outgoing interfaces that have the H-bit flag are switched in hardware and the remaining traffic is switched in software through the RP. For this reason, it may be desirable to rate limit the flow destined to the RP for forwarding and replication, which might otherwise increase CPU utilization.

The multicast directly connected rate limiter limits the multicast packets from directly connected sources.

This example shows how to rate limit the multicast packets to 30000 pps with a burst of 30:

```
Router(config)# mls rate-limit multicast ipv4 connected 30000 30
```

The **ip-option** keyword and the ip-option rate limiter are not supported in PFC3A mode.

This example shows how to set the rate limiters for the IPv4 multicast packets failing the uRPF check:

```
Router(config)# mls rate-limit multicast ipv4 non-rpf 100
```

This example shows how to rate limit the multicast FIB miss packets to 10000 pps with a burst of 10:

```
Router(config)# mls rate-limit multicast ipv4 fib-miss 10000 10
```

This example shows how to rate limit the partial shortcut flows to 20000 pps with a burst of 20 packets:

```
Router(config)# mls rate-limit multicast ipv4 partial 20000 20
```

This example shows how to rate limit the multicast packets to 30000 pps with a burst of 20:

```
Router(config)# mls rate-limit multicast ipv4 connected 30000 20
```

This example shows how to rate limit IGMP-snooping traffic:

```
Router(config)# mls rate-limit multicast ipv4 igmp 20000 40
```

IPv6 Multicast

This rate limiter limits the IPv6 multicast packets. Table 52-1 lists the IPv6 rate limiters and the class of traffic that each rate limiter serves.

Table 52-1   IPv6 Rate Limiters

Rate Limiter	Traffic Classes to be Rate Limited
Connected	Directly connected source traffic
Default-drop	* (*, G/m) SSM
	* (*, G/m) SSM non-rpf
Route-control	* (*, FF02::X/128)
Starg-bridge	* (*, G/128) SM
	* SM non-rpf traffic when (*, G) exists
Starg-M-bridge	* (*, G/m) SM
	* (*, FF/8)
	* SM non-rpf traffic when (*, G) doesn't exist

You can configure rate limiters for IPv6 multicast traffic using one of the following methods:

- Direct association of the rate limiters for a traffic class—Select a rate and associate the rate with a rate limiter. This example shows how to pick a rate of 1000 pps and 20 packets per burst and associate the rate with the **default-drop** rate limiter:

  ```
  Router(config)# mls rate-limit multicast ipv6 default-drop 1000 20
  ```

- Static sharing of a rate limiter with another preconfigured rate limiter—When there are not enough adjacency-based rate limiters available, you can share a rate limiter with an already configured rate limiter (target rate limiter). This example shows how to share the **route-cntl** rate limiter with the **default-drop** target rate limiter:

  ```
  Router(config)# mls rate-limit multicast ipv6 route-cntl share default-drop
  ```

  If the target rate limiter is not configured, a message is displayed that indicates that the target rate limiter must be configured for it to be shared with other rate limiters.

- Dynamic sharing of rate limiters—If you are not sure about which rate limiter to share with, use the **share auto** keywords to enable dynamic sharing. When you enable dynamic sharing, the system selects a preconfigured rate limiter and shares the given rate limiter with the preconfigured rate limiter. This example shows how to choose dynamic sharing for the route-cntrl rate limiter:

  ```
  Router(config)# mls rate-limit multicast ipv6 route-cntl share auto
  ```

This example shows how to set the rate limiters for the IPv6 multicast packets from a directly connected source:

```
Router(config)# mls rate-limit multicast ipv6 connected 1500 20
```

This example shows how to configure a direct association of the rate limiters for a traffic class:

```
Router(config)# mls rate-limit multicast ipv6 default-drop 1000 20
```

This example shows how to configure the static sharing of a rate limiter with another preconfigured rate limiter:

```
Router(config)# mls rate-limit multicast ipv6 route-cntl share default-drop
```

This example shows how to enable dynamic sharing for the route control rate limiter:

```
Router(config)# mls rate-limit multicast ipv6 route-cntl share auto
```

DoS Protection Default Configuration

Table 52-2 shows the DoS protection default configuration for the PFC3 hardware-based rate limiters.

Table 52-2   PFC3 Hardware-based Rate Limiter Default Setting

Rate Limiter	Default Status (ON/OFF)	Default Value
Ingress/Egress ACL Bridged Packets	OFF	
RPF Failures	ON	100 pps, burst of 10 packets
FIB Receive cases	OFF	
FIB Glean Cases	OFF	
Layer 3 Security features	OFF	
ICMP Redirect	OFF	
ICMP Unreachable	ON	100 pps, burst of 10 packets
VACL Log	ON	2000 pps, burst of 10 packets
TTL Failure	OFF	
MTU Failure	OFF	
Layer 2 PDU	OFF	
Layer 2 Protocol Tunneling	OFF	
IP Errors	ON	100 pps, burst of 10 packets
Multicast IGMP	OFF	
Multicast FIB-Miss	ON	100000 pps, burst of 100 packets
Multicast Partial-SC	ON	100000 pps, burst of 100 packets
Multicast Directly Connected	OFF	
Multicast Non-RPF	OFF	
Multicast IPv6	ON	If the *packets-in-burst* is not set, a default of **100** is programmed for multicast cases.

DoS Protection Configuration Guidelines and Restrictions

When configuring DoS protection, follow these CPU rate limiter guidelines and restrictions:

Note   For the CoPP guidelines and restrictions, see the "CoPP Configuration Guidelines and Restrictions" section on page 53-2.

- Do not use these rate limiters if multicast is enabled in systems configured with a PFC3A:
    - TTL failure
    - MTU failure
- These rate limiters are not supported in PFC3A mode:
    - Unicast IP options
    - Multicast IP options
- These are Layer 2 rate limiters:
    - Layer 2 PDUs
    - Layer 2 protocol tunneling
    - Layer 2 Multicast IGMP
- There are eight Layer 3 registers and two Layer 2 registers that can be used as CPU rate limiters.
- Do not use the CEF receive limiter if CoPP is being used. The CEF receive limiter will override the CoPP traffic.
- Rate limiters override the CoPP traffic.
- Configured rate limits is applied to each forwarding engine (except for the Layer 2 hardware rate limiter which is applied globally).
- Layer 2 rate limiters are not supported in truncated mode.
- The following restrictions apply when using the ingress and egress ACL-bridged packet rate limiters:
    - The ingress and egress ACL-bridged packet rate limiter is available for unicast traffic only.
    - The ingress and egress ACL-bridged packet rate limiters share a single rate-limiter register. If you enable the ACL-bridge ingress and egress rate limiters, both the ingress and the egress ACLs must share the same rate-limiter value.
- Use the **mls rate-limit unicast** command to rate limit unicast traffic.
- Use the **mls rate-limit multicast** command to rate limit multicast traffic.
- Use the **mls rate-limit multicast layer 2** command to rate limit Layer 2 multicast traffic.

Monitoring Packet Drop Statistics

You can capture the incoming or outgoing traffic on an interface and send a copy of this traffic to an external interface for monitoring by a traffic analyzer. To capture traffic and forward it to an external interface, use the **monitor session** command.

When capturing traffic, these restrictions apply:

- The incoming captured traffic is not filtered.
- The incoming captured traffic is not rate limited to the capture destination.

Monitoring Dropped Packets Using Monitor Session Commands

This example shows how to use the **monitor session** command to capture and forward traffic to an external interface:

```
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
```

```
Router(config)# monitor session 1 source vlan 44 both
Router(config)# monitor session 1 destination interface g9/1
Router(config)# end
Router#
2w0d: %SYS-5-CONFIG_I: Configured from console by console
```

This example shows how to use the **show monitor session** command to display the destination port location:

```
Router# show monitor session 1
Session 1
---------
Source Ports:
    RX Only:        None
    TX Only:        None
    Both:           None
Source VLANs:
    RX Only:        None
    TX Only:        None
    Both:           44
Destination Ports: Gi9/1
Filter VLANs:       None
```

Monitoring Dropped Packets Using show tcam interface Command

All modes except PFC3A mode support ACL hit counters in hardware. You can use the **show tcam interface** command to display each entry in the ACL TCAM.

This example shows how to use the **show tcam interface** command to display the number of times the entry was hit:

```
Router# show tcam interface fa5/2 acl in ip detail

-----------------------------------------------------------------------------------------------------
DPort - Destination Port   SPort - Source Port      TCP-F - U -URG Pro  - Protocol
I     - Inverted LOU       TOS   - TOS Value              - A -ACK rtr  - Router
MRFM  - M -MPLS Packet      TN    - T -Tcp Control        - P -PSH COD  - C -Bank Care Flag
      - R -Recirc. Flag           - N -Non-cachable       - R -RST      - I -OrdIndep. Flag
      - F -Fragment Flag    CAP   - Capture Flag          - S -SYN      - D -Dynamic Flag
      - M -More Fragments   F-P   - FlowMask-Prior.       - F -FIN T    - V(Value)/M(Mask)/R(Result)
X     - XTAG                (*)   - Bank Priority
-----------------------------------------------------------------------------------------------------


Interface: 1018   label: 1   lookup_type: 0
protocol: IP   packet-type: 0


+-+-----+--------------+--------------+--------------+--------------+------+---+----+-+---+--+---+---+
|T|Index|  Dest Ip Addr | Source Ip Addr|    DPort     |    SPort     | TCP-F|Pro|MRFM|X|TOS|TN|COD|F-P|
+-+-----+--------------+--------------+--------------+--------------+------+---+----+-+---+--+---+---+
 V 18396      0.0.0.0        0.0.0.0        P=0            P=0        ------  0 ---- 0   0 -- --- 0-0
 M 18404      0.0.0.0        0.0.0.0         0              0                 0 ---- 0   0
 R rslt: L3_DENY_RESULT            rtr_rslt: L3_DENY_RESULT


 V 36828      0.0.0.0        0.0.0.0        P=0            P=0        ------  0 ---- 0   0 -- --- 0-0
 M 36836      0.0.0.0        0.0.0.0         0              0                 0 ---- 0   0
 R rslt: L3_DENY_RESULT (*)        rtr_rslt: L3_DENY_RESULT (*)
Router#
```

You can also use the TTL and IP options counters to monitor the performance of the Layer 3 forwarding engine.

This example shows how to use the **show mls statistics** command to display packet statistics and errors associated with the Layer 3 forwarding engine:

```
Router# show mls statistics

Statistics for Earl in Module 6

L2 Forwarding Engine
  Total packets Switched              : 25583421

L3 Forwarding Engine
  Total packets L3 Switched           : 25433414 @ 24 pps

  Total Packets Bridged               : 937860
  Total Packets FIB Switched          : 23287640
  Total Packets ACL Routed            : 0
  Total Packets Netflow Switched      : 0
  Total Mcast Packets Switched/Routed : 96727
  Total ip packets with TOS changed   : 2
  Total ip packets with COS changed   : 2
  Total non ip packets COS changed    : 0
  Total packets dropped by ACL        : 33
  Total packets dropped by Policing   : 0

Errors
  MAC/IP length inconsistencies       : 0
  Short IP packets received           : 0
  IP header checksum errors           : 0
  TTL failures                        : 0
<---------------- TTL counters
  MTU failures                        : 0
<----------------MTU failure counters

Total packets L3 Switched by all Modules: 25433414 @ 24 pps
```

Monitoring Dropped Packets Using VACL Capture

The VACL capture feature allows you to direct traffic to ports configured to forward captured traffic. The capture action sets the capture bit for the forwarded packets so that ports with the capture function enabled can receive the packets. Only forwarded packets can be captured.

You can use VACL capture to assign traffic from each VLAN to a different interface.

VACL capture does not allow you to send one type of traffic, such as HTTP, to one interface and another type of traffic, such as DNS, to another interface. Also, VACL capture granularity is only applicable to traffic switched locally; you cannot preserve the granularity if you direct traffic to a remote switch.

This example shows how to use VACL capture to capture and forward traffic to a local interface:

```
Router(config-if)# switchport capture
Router(config-if)# switchport capture allowed vlan add 100
```

Displaying Rate-Limiter Information

The **show mls rate-limit** command displays information about the configured rate limiters.

The **show mls rate-limit usage** command displays the hardware register that is used by a rate-limiter type. If the register is not used by any rate-limiter type, Free is displayed in the output. If the register is used by a rate-limiter type, Used and the rate-limiter type are displayed.

In the command output, the rate-limit status could be one of the following:

- On indicates that a rate for that particular case has been set.

- Off indicates that the rate-limiter type has not been configured, and the packets for that case are not rate limited.

- On/Sharing indicates that a particular case (not manually configured) is affected by the configuration of another rate limiter belonging to the same sharing group.

- A hyphen indicates that the multicast partial-SC rate limiter is disabled.

In the command output, the rate-limit sharing indicates the following information:

- Whether sharing is static or dynamic

- Group dynamic sharing codes

To display the configured rate limiters, use the **show mls rate-limit** command:

```
Router# show mls rate-limit
 Sharing Codes: S - static, D - dynamic
 Codes dynamic sharing: H - owner (head) of the group, g - guest of the group

   Rate Limiter Type      Status    Packets/s   Burst  Sharing
 --------------------    ----------  ---------   -----  -------
         MCAST NON RPF    Off           -          -      -
        MCAST DFLT ADJ    On          100000      100   Not sharing
      MCAST DIRECT CON    Off           -          -      -
        ACL BRIDGED IN    Off           -          -      -
       ACL BRIDGED OUT    Off           -          -      -
           IP FEATURES    Off           -          -      -
          ACL VACL LOG    On           2000        1    Not sharing
           CEF RECEIVE    Off           -          -      -
             CEF GLEAN    Off           -          -      -
       MCAST PARTIAL SC   On          100000      100   Not sharing
       IP RPF FAILURE     On           100        10    Group:0 S
          TTL FAILURE     Off           -          -      -
 ICMP UNREAC. NO-ROUTE    On           100        10    Group:0 S
 ICMP UNREAC. ACL-DROP    On           100        10    Group:0 S
         ICMP REDIRECT    Off           -          -      -
           MTU FAILURE    Off           -          -      -
        MCAST IP OPTION   Off           -          -      -
        UCAST IP OPTION   Off           -          -      -
          LAYER_2 PDU     Off           -          -      -
           LAYER_2 PT     Off           -          -      -
            IP ERRORS     On           100        10    Group:0 S
           CAPTURE PKT    Off           -          -      -
            MCAST IGMP    Off           -          -      -
 MCAST IPv6 DIRECT CON    Off           -          -      -
 MCAST IPv6 *G M BRIDG    Off           -          -      -
  MCAST IPv6 *G BRIDGE    Off           -          -      -
  MCAST IPv6 SG BRIDGE    Off           -          -      -
 MCAST IPv6 ROUTE CNTL    Off           -          -      -
 MCAST IPv6 DFLT DROP     Off           -          -      -
 MCAST IPv6 SECOND. DR    Off           -          -      -
Router#
```

To display the usage of the hardware rate limiters, use the **show mls rate-limit usage** command:

```
Router# show mls rate-limit usage
                        Rate Limiter Type    Packets/s   Burst
                       --------------------   ---------   -----
Layer3 Rate Limiters:
          RL# 0: Free                            -          -      -
          RL# 1: Free                            -          -      -
```

```
                              RL# 2: Free                   -          -        -
                              RL# 3: Used
                                              MCAST DFLT ADJ    100000      100
                              RL# 4: Free                   -          -        -
                              RL# 5: Free                   -          -        -
                              RL# 6: Used
                                                IP RPF FAILURE       100       10
                                     ICMP UNREAC. NO-ROUTE            100       10
                                     ICMP UNREAC. ACL-DROP            100       10
                                                    IP ERRORS        100       10
                              RL# 7: Used
                                                ACL VACL LOG        2000        1
                              RL# 8: Rsvd for capture        -          -        -

                  Layer2 Rate Limiters:
                              RL# 9: Reserved
                              RL#10: Reserved
                              RL#11: Free                   -          -        -
                              RL#12: Free                   -          -        -
                  Router#
```

Configuring Sticky ARP

Sticky ARP prevents MAC address spoofing by ensuring that ARP entries (IP address, MAC address, and source VLAN) do not get overridden. The switch maintains ARP entries in order to forward traffic to end devices or other switches. ARP entries are usually updated periodically or modified when ARP broadcasts are received. During an attack, ARP broadcasts are sent using a spoofed MAC address (with a legitimate IP address) so that the switch learns the legitimate IP address with the spoofed MAC address and begins to forward traffic to that MAC address. With sticky ARP enabled, the switch learns the ARP entries and does not accept modifications received through ARP broadcasts. If you attempt to override the sticky ARP configuration, you will receive an error message. For a complete description of the system error messages, see the *Catalyst 6500 Series Switch Cisco IOS System Message Guide*, Release 12.2SX at this URL:

http://www.cisco.com/en/US/docs/ios/12_2sx/system/messages/122sxsms.html

To configure sticky ARP on a Layer 3 interface, perform this task:

	Command	Purpose
Step 1	Router(config)# **interface** *type*[1] *slot/port*	Selects the interface on which sticky ARP is applied.
Step 2	Router(config-if)# **ip sticky-arp**	Enables sticky ARP.
	Router(config-if)# **no ip sticky-arp ignore**	Removes the previously configured sticky ARP command.
Step 3	Router(config-if)# **ip sticky-arp ignore**	Disables sticky ARP.

1. *type* = **fastethernet**, **gigabitethernet**, or **tengigabitethernet**

This example shows how to enable sticky ARP on interface 5/1:

```
Router# configure terminal
Router(config)# interface gigabitethernet 5/1
Router(config-if)# ip sticky-arp
Router(config-if)# end
Router#
```

Tip   For additional information about Cisco Catalyst 6500 Series Switches (including configuration examples and troubleshooting information), see the documents listed on this page:

http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_series_home.html

Participate in the Technical Documentation Ideas forum