Denial of Service (DoS) Protection

Note

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

  http://www.cisco.com/en/US/products/ps11846/prod_command_reference_list.html
- Cisco IOS Release 15.1SY supports only Ethernet interfaces. Cisco IOS Release 15.1SY does not support any WAN features or commands.
- Also see:
  - Chapter 1, “MAC Address-Based Traffic Blocking”
  - Chapter 1, “Traffic Storm Control”
  - Chapter 1, “Control Plane Policing (CoPP)”
  - http://www.cisco.com/en/US/docs/ios-xml/ios/security/config_library/15-sy/secdata-15-sy-library.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

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.

See Chapter 1, “Cisco IOS ACL Support,” and Chapter 1, “VLAN ACLs (VACLs).”

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 PFC and DFCs perform 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
```

See Chapter 1, "PFC QoS."

Global Protocol Packet Policing

Prerequisites for Global Protocol Packet Policing

None.

Restrictions for Global Protocol Packet Policing

- The minimum values supported by the **mls qos protocol arp police** command are too small for use in production networks.
- 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.
- 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.
- Policy-based protocol packet policing is applied per-forwarding engine (PFC and any DFCs).
- 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.
- Policy-based protocol packet policing does not support microflow policers.
- Only ingress Policy-based protocol packet policing is supported.
- Policy-based protocol packet policing does not support Layer 4 ACL operators (see the "Restrictions for Layer 4 Operators in ACLs" section on page 1-2), which imposes these subsequent restrictions:
  - For IPv4 or IPv6 traffic, no support for UDP or TCP port range matching
  - For IPv6 traffic, no support for precedence or DSCP matching
- Protocol packet policing policies can share an aggregate policer with QoS policies.
- An aggregate policer cannot be applied to both ingress and egress traffic.

- Policy-based protocol packet policing supports the **class default** and **permit** *protocol_name* **any any** commands, but traffic flow might be affected significantly because the protocol packet policing policy processes all matched traffic.
- With Supervisor Engine 720, policy-based protocol packet policing is applied only to untrusted ports.
- You can configure both single-command protocol packet policing and policy-based protocol packet policing. Single-command protocol packet policing is applied first and then policy-based protocol packet policing.

Note The software does not detect or attempt to resolve any configuration conflicts between single-command protocol packet policing and policy-based protocol packet policing.

- You can configure both policy-based protocol packet policing and control plane policing (see Chapter 1, "Control Plane Policing (CoPP)"). Policy-based protocol packet policing is applied first and then CoPP.
- Single-command protocol packet policing programs the configured protocol-specific action for ingress traffic and automatically programs a corresponding egress-traffic pass-through action to preserve the ingress result egress traffic.
- Policy-based protocol packet policing does not automatically preserve the ingress policing result in egress traffic.
  - To preserve the ingress policing result in egress traffic with policy-based protocol packet policing, configure an appropriate output policy. To pass egress traffic through unchanged, duplicate each ingess class in the output policy and configure **trust dscp** as the class-map action.
  - Without an output policy-map, egress traffic is processed by any configured interface-based policy-map and ingress global policy result will be overwritten.
- The PFC and any DFCs supports a single **match** command in **class-map match-all** class maps, except that the **match protocol** command can be configured in a class map with the **match dscp** or **match precedence** command.
- The PFC and any DFCs supports multiple **match** commands in **class-map match-any** class maps.
- Class maps can use the **match** commands listed in Table 1-1 to configure a traffic class that is based on the match criteria.

Table 1-1 Traffic Classification Class Map match Commands and Match Criteria

match Commands	Direction	Match Criteria
match access-group {*access_list_number* | **name** *access_list_name*}	Ingress	Access control list (ACL). **Note** Use ACLs to match the following: —CoS value —VLAN ID —Packet length
match any	Ingress	Any match criteria.
match cos	Ingress	CoS value.
match discard-class	Ingress	Discard class value.

Table 1-1 Traffic Classification Class Map match Commands and Match Criteria (continued)

match Commands	Direction	Match Criteria
match dscp **Note** The **match protocol** command can be configured in a class map with the **match dscp** command.	Ingress	DSCP value.
match l2 miss	Ingress	Layer 2 traffic flooded in a VLAN because it is addressed to a currently unlearned MAC-Layer destination address.
match mpls experimental topmost	Ingress	MPLS EXP value in the topmost label.
match precedence **Note** The **match protocol** command can be configured in a class map with the **match precedence** command.	Ingress	IP precedence values.
match protocol {**arp** | **ip** | **ipv6**} **Note** The **match protocol** command can be configured in a class map with the **match dscp** or **match precedence** command.	Ingress	Protocol.
match qos-group	Ingress	QoS group ID.

The PFC and any DFCs supports these ACL types for use with the **match access group** command:

Protocol	Numbered ACLs	Extended ACLs	Named ACLs
IPv4	Yes: 1 to 99 1300 to 1999	Yes: 100 to 199 2000 to 2699	Yes
IPv6	Not applicable	Yes (named)	Yes
MAC Layer	Not applicable	Not applicable	Yes
ARP	Not applicable	Not applicable	Yes

Information About Global Protocol Packet Policing

Attackers may try to overwhelm the RP CPU with routing protocol control packets (for example, ARP packets). Protocol packet policing rate limits this traffic in hardware. Release 15.1(1)SY1 and later releases support policy-based global protocol packet policing, shown in Cisco Feature Navigator as the Global QoS Policy feature.

How to Configure Single-Command Global Protocol Packet Policing

Enter the **mls qos protocol ?** to display the supported routing protocols.

The **mls qos protocol arp police** command rate limits ARP packets. This example shows how to allow 200 ARP requests and replies per second:

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

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)
```

How to Configure Policy-Based Global Protocol Packet Policing

Use these QoS sections and the global protocol packet policing policy map configuration section:

- Configuring a Class Map, page 1-70
- Configuring a Policy Map, page 1-72
- Configuring a Global Protocol Packet Policing Policy Map, page 1-6

Configuring a Global Protocol Packet Policing Policy Map

To configure a global protocol packet policing policy map, perform this task:

Command	Purpose
`Router(config)# mls qos service-policy input policy_map_name`	Configures a global protocol packet policing policy map. **Note** You can configure one input policy.

Unicast Reverse Path Forwarding (uRPF) Check

- Prerequisites for uRPF Check, page 1-7
- Restrictions for uRPF Check, page 1-7
- Information about uRPF Check, page 1-8
- Configuring the Unicast RPF Check Mode, page 1-9
- Enabling Self-Pinging, page 1-11

Prerequisites for uRPF Check

None.

Restrictions for uRPF Check

- Unicast RPF does not provide complete protection against spoofing. Spoofed packets can enter a network through unicast RPF-enabled interfaces if an appropriate return route to the source IP address exists.
- Avoid configurations that overload the route processor with unicast RPF checks.
  - Do not configure unicast RPF to filter with an ACL.
  - Do not configure the global unicast RPF “punt” check mode.
- The PFC does not provide hardware support for the unicast RPF check for policy-based routing (PBR) traffic. (CSCea53554)
- The switch applies the same unicast RPF mode to all interfaces where unicast RPF check is configured. Any change that you make in the unicast RPF mode on any interfaces is applied to all interfaces where the unicast RPF check is configured.
- The “allow default” options of the unicast RPF modes do not offer significant protection against spoofing.
  - Strict unicast RPF Check with Allow Default—Received IP traffic that is sourced from a prefix that exists in the routing table passes the unicast RPF check if the prefix is reachable through the input interface. If a default route is configured, any IP packet with a source prefix that is not in the routing table passes the unicast RPF check if the ingress interface is a reverse path for the default route.
  - Loose unicast RPF Check with Allow Default—If a default route is configured, any IP packet passes the unicast RPF check.
  - If, on a maximum of 24 interfaces, divided into four groups of six interfaces each, the switch receives valid IP packets that have up to six reverse-path interfaces per source prefix, you can configure the unicast RPF strict mode with the **mls ip cef rpf multipath interface-group** command.

    This option requires you to identify the source prefixes and the interfaces that serve as reverse paths for the source prefixes and to configure interface groups for those reverse path interfaces. All of the reverse-path interfaces for each source prefix must be in the same interface group. You can configure four interface groups, with each group containing up to six reverse-path interfaces. There is no limit on the number of source prefixes that an interface group can support.

    To ensure that no more than six reverse-path interfaces exist in the routing table for each prefix, enter the **maximum-paths 6** command in config-router mode when configuring OSPF, EIGRP, or BGP.

    IP traffic with one or two reverse-path interfaces that is received on uPPF-check enabled interfaces outside the interface groups passes the unicast RPF check if the ingress interface and at most one other interface are reverse paths.

    With maximum paths set to six, IP traffic with more than two reverse-path interfaces that is received on uPPF-check enabled interfaces outside the interface groups always pass the unicast RPF check.

  - If, on any number of interfaces, the switch receives valid IP packets that have one or two reverse path interfaces per source prefix, you can configure the unicast RPF strict mode with the **mls ip cef rpf multipath pass** command.

    To ensure that no more than two reverse-path interfaces exist in the routing table for each prefix, enter the **maximum-paths 2** command in config-router mode when configuring OSPF, EIGRP, or BGP.

- Unicast RPF Loose Mode with the Pass Global Mode—The unicast RPF loose mode provides less protection than strict mode, but it is an option on switches that receive valid IP traffic on interfaces that are not reverse paths for the traffic. The unicast RPF loose mode verifies that received traffic is sourced from a prefix that exists in the routing table, regardless of the interface on which the traffic arrives.

Information about uRPF Check

The unicast RPF check verifies that the source address of received IP packets is reachable. The unicast RPF check discards IP packets that lack a verifiable IP source prefix (route), which helps mitigate problems that are caused by traffic with malformed or forged (spoofed) IP source addresses.

The unicast RPF check is performed in software on the RP and supports up to 16 reverse-path interfaces.

To ensure that no more than 16 reverse-path interfaces exist in the routing table for each prefix, enter the **maximum-paths 16** command in config-router mode when configuring OSPF, EIGRP, or BGP.

For unicast RPF check without ACL filtering, the PFC3 provides hardware support for the RPF check of traffic from multiple interfaces.

For unicast RPF check with ACL filtering, the PFC determines whether or not traffic matches the ACL. The PFC sends the traffic denied by the RPF ACL to the route processor (RP) for the unicast RPF check. Packets permitted by the ACL are forwarded in hardware without a unicast RPF check.

How to Configure Unicast RPF Check

Configuring the Unicast RPF Check Mode

To configure unicast RPF check mode, perform this task:

	Command	Purpose
Step 1	`Router(config)# interface {{vlan vlan_ID} | {type slot/port} | {port-channel number}}`	Selects an interface to configure. **Note** Based on the input port, unicast RPF check verifies the best return path before forwarding the packet on to the next destination.
Step 2	`Router(config-if)# ip verify unicast source reachable-via {rx | any} [allow-default] [list]`	Configures the IPv4 unicast RPF check mode.
Step 3	`Router(config-if)# ipv6 verify unicast source reachable-via {rx | any} [allow-default] [list]`	Configures the IPv6 unicast RPF check mode.
Step 4	`Router(config-if)# exit`	Exits interface configuration mode.
Step 5	`Router# show mls hardware cef ip rpf`	Verifies the IPv4 configuration.
Step 6	`Router# show platform hardware cef ipv6 rpf`	Verifies the IPv6 configuration.

Note The most recently configured mode is automatically applied to all ports configured for unicast RPF check.

- Use the **rx** keyword to enable strict check mode.
- Use the **any** keyword to enable exist-only check mode.
- Use the **allow-default** keyword to allow use of the default route for RPF verification.
- Use the *list* option to identify an access list.
  - If the access list denies network access, denied packets are dropped at the port.
  - If the access list permits network access, packets are forwarded to the destination address. Forwarded packets are counted in the interface statistics.
  - If the access list includes the logging action, information about the packets is sent to the log server.

This example shows how to enable unicast RPF exist-only check mode on Gigabit Ethernet port 4/1:

```
Router(config)# interface gigabitethernet 4/1
Router(config-if)# ipv6 verify unicast source reachable-via any
Router(config-if)# ip verify unicast source reachable-via any
Router(config-if)# end
Router#
```

This example shows how to enable unicast RPF strict check mode on Gigabit Ethernet port 4/2:

```
Router(config)# interface gigabitethernet 4/2
Router(config-if)# ipv6 verify unicast source reachable-via rx
Router(config-if)# ip verify unicast source reachable-via rx
Router(config-if)# end
Router#
```

Note ACL-based unicast RPF check is not supported in the hardware on Supervisor Engine 720. If you configure ACL-based uRPF check on an interface, all packets denied in the ACL are redirected to the MSFC3 CPU for the uRPF check, while packets permitted by the ACL are forwarded in hardware without a uRPF check. Redirected packets are subject to the global ACL bridged input hardware-to-CPU rate limiter (that limits the amount of such traffic that reaches the MSFC3 CPU) and may be dropped in hardware. Configure this rate limiter with the **mls rate-limit unicast acl input** global configuration command.

Beginning in Cisco IOS release 12.2(18)SXF6, you can change this behavior using the **mls ip cef rpf hw-enable-rpf-acl** global configuration command. When you add this command, packets permitted by the exception ACL are redirected to the MSFC3 CPU for the uRPF check, while packets denied by the ACL are forwarded in hardware with uRPF check. If the packets fail the hardware RPF check, they are punted to CPU. RPF-fail packets are subject to the global RPF fail hardware-to-CPU rate limiter. Configure this rate limiter with the **mls rate-limit unicast ip rpf-failure** global configuration command.

Configuring the Multiple-Path Unicast RPF Check Mode

To configure the multiple-path unicast RPF check mode, perform this task:

	Command	Purpose
Step 1	`Router(config)# mls ip cef rpf mpath {punt | pass | interface-group}`	Configures the multiple path RPF check mode.
Step 2	`Router(config)# end`	Exits configuration mode.
Step 3	`Router# show mls cef ip rpf`	Verifies the configuration.

When configuring multiple path RPF check, note the following information:

- **punt** mode (default)—The PFC3 performs the unicast RPF check in hardware for up to two interfaces per prefix. Packets arriving on any additional interfaces are redirected (punted) to the RP for unicast RPF check in software.
- **pass** mode—The PFC3 performs the unicast RPF check in hardware for single-path and two-path prefixes. Unicast RPF check is disabled for packets coming from multipath prefixes with three or more reverse-path interfaces (these packets always pass the unicast RPF check).
- **interface-group** mode—The PFC3 performs the unicast RPF check in hardware for single-path and two-path prefixes. The PFC3 also performs the unicast RPF check for up to four additional interfaces per prefix through user-configured multipath unicast RPF check interface groups. Unicast RPF check is disabled for packets coming from other multipath prefixes that have three or more reverse-path interfaces (these packets always pass the unicast RPF check).

This example shows how to configure punt as the multiple path RPF check mode:

```
Router(config)# mls ip cef rpf mpath punt
```

Configuring Multiple-Path Interface Groups

To configure multiple-path unicast RPF interface groups, perform this task:

	Command	Purpose
Step 1	`Router(config)# mls ip cef rpf interface-group [0 | 1 | 2 | 3] interface1 [interface2 [interface3 [interface4]]]`	Configures a multiple path RPF interface group.
Step 2	`Router(config)# mls ip cef rpf interface-group group_number`	Removes an interface group.
Step 3	`Router(config)# end`	Exits configuration mode.
Step 4	`Router# show mls cef ip rpf`	Verifies the configuration.

This example shows how to configure interface group 2:

```
Router(config)# mls ip cef rpf interface-group 2 fastethernet 3/3 fastethernet 3/4
fastethernet 3/5 fastethernet 3/6
```

Enabling Self-Pinging

With unicast RPF check enabled, by default the switch cannot ping itself. To enable self-pinging, perform this task:

	Command	Purpose
Step 1	`Router(config)# interface {{vlan vlan_ID} | {type slot/port} | {port-channel number}}`	Selects the interface to configure.
Step 2	`Router(config-if)# ip verify unicast source reachable-via any allow-self-ping`	Enables the switch to ping itself or a secondary address.
Step 3	`Router(config-if)# exit`	Exits interface configuration mode.

This example shows how to enable self-pinging:

```
Router(config)# interface gigabitethernet 4/1
Router(config-if)# ip verify unicast source reachable-via any allow-self-ping
Router(config-if)# end
```

Hardware-Based Rate Limiters

Prerequisites for Rate-Limiters

None.

Restrictions for Rate-Limiters

- These are Layer 2 rate limiters:
  - Layer 2 PDUs
  - Layer 2 protocol tunneling
  - Layer 2 Multicast IGMP
- Do not use the CEF receive limiter if CoPP is being used. The CEF receive limiter will 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.

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.

Ingress-Egress ACL Bridged Packets (Unicast Only)

Commands:

mls rate-limit unicast acl input

mls rate-limit unicast acl output

The PFC and DFC provide separate ACL-bridge rate-limiters.

This rate limiter rate limits packets sent to the RP because of an ingress or egress ACL bridge results. 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. If the ACL bridge rate limiting is disabled, the Layer 3 redirect rate limit results are converted to the bridge result.

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.

uRPF Check Failure

Command: **mls rate-limit unicast ip rpf-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.

TTL Failure

Command: **mls rate-limit all 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.

ICMP Unreachable (Unicast Only)

Commands:

mls rate-limit unicast ip icmp unreachable acl-drop

mls rate-limit unicast ip icmp unreachable no-route

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.

FIB (CEF) Receive Cases (Unicast Only)

Command: **mls rate-limit unicast cef receive**

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.

FIB Glean (Unicast Only)

Command: **mls rate-limit unicast cef glean**

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.

Layer 3 Security Features (Unicast Only)

Command: **mls rate-limit unicast ip features**

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.

ICMP Redirect (Unicast Only)

Command: **mls rate-limit unicast ip icmp redirect**

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.

VACL Log (Unicast Only)

Command: **mls rate-limit unicast acl vacl_log**

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.

MTU Failure

Command: **mls rate-limit all mtu**

Like 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.

Layer 2 Multicast IGMP Snooping

Command: **mls rate-limit multicast ipv4 igmp**

The IGMP snooping rate limiter limits the number of Layer 2 IGMP packets destined for the RP. 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. The IGMP snooping rate limiter does not rate limit PIM messages.

Layer 2 PDU

Command: **mls rate-limit layer2 pdu**

The Layer 2 protocol data unit (PDU) rate limiter allows you to limit the number of hardware-switched Layer 2 PDU protocol packets (including BPDUs, DTP, PAgP, CDP, STP, and VTP packets). 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.

Layer 2 Protocol Tunneling

Command: **mls rate-limit layer2 l2pt**

This rate limiter limits the hardware-switched Layer 2 protocol tunneling packets, which include control PDUs, CDP, STP, and VTP packets. 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.

IP Errors

Command: **mls rate-limit unicast ip errors**

This rate limiter limits the packets with IP checksum and length errors. When a packet reaches the PFC or DFC 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.

IPv4 Multicast

Commands:

mls rate-limit multicast ipv4 connected

mls rate-limit multicast ipv4 fib-miss

mls rate-limit multicast ipv4 igmp

mls rate-limit multicast ipv4 ip-options

mls rate-limit multicast ipv4 pim

These rate limiters limit 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.

IPv6 Multicast

Commands:

mls rate-limit multicast ipv6 connected

mls rate-limit multicast ipv6 control-packet

mls rate-limit multicast ipv6 mld

These rate limiters limit IPv6 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.

Hardware-Based Rate Limiters Default Configuration

Table 1-1 shows the DoS protection default configuration for the hardware-based rate limiters.

Table 1-1 Hardware-based Rate Limiter Default Settings

Rate Limiter	Default State	Default Value
CEF RECEIVE	Off	
CEF RECEIVE SECONDARY	On	pps: 15000; burst microseconds: 1000000
CEF GLEAN	On	pps: 1000; burst microseconds:1000000
IP ERRORS	Off	
UCAST IP OPTION	On	pps: 100; burst microseconds:1000000
ICMP ACL-DROP	On	pps: 100; burst microseconds:1000000
ICMP NO-ROUTE	On	pps: 100; burst microseconds: 1000000
ICMP REDIRECT	Off	
RPF FAILURE	On	pps: 100; burst microseconds: 1000000
ACL VACL LOG	On	pps: 2000; burst microseconds: 1000000
ACL BRIDGED IN	Off	
ACL BRIDGED OUT	Off	
ARP Inspection	Off	
DHCP Snooping IN	Off	
IP FEATURES	Off	
MAC PBF IN	Off	
CAPTURE PKT	Off	
IP ADMIS. ON L2 PORT	Off	
MCAST IPV4 DIRECTLY C	Off	
MCAST IPV4 FIB MISS	Off	
MCAST IPV4 IGMP	Off	
MCAST IPV4 OPTIONS	Off	
MCAST IPV4 PIM	Off	
MCAST IPV6 DIRECTLY C	Off	
MCAST IPV6 MLD	Off	
MCAST IPV6 CONTROL PK	Off	
MTU FAILURE	Off	
TTL FAILURE	Off	
MCAST BRG FLD IP CNTR	Off	

Table 1-1 Hardware-based Rate Limiter Default Settings (continued)

Rate Limiter	Default State	Default Value
MCAST BRG FLD IP	Off	
MCAST BRG	Off	
MCAST BRG OMF	Off	
UCAST UNKNOWN FLOOD	Off	
LAYER_2 PDU	Off	
LAYER_2 PT	Off	
LAYER_2 PORTSEC	Off	
LAYER_2 SPAN PCAP	Off	
DIAG RESERVED 0	On	pps: 33554431; burst microseconds: 1
DIAG RESERVED 1	On	pps: 33554431; burst microseconds: 1
DIAG RESERVED 2	On	pps: 33554431; burst microseconds: 1
DIAG RESERVED LIF 0	On	pps: 33554431; burst microseconds: 1
MCAST REPL RESERVED	On	pps: 1; burst microseconds: 0

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
 State : ON - enabled but not sharing, ON/S - enabled and sharing
 Share : NS - not sharing, G - group, S - static sharing, D - dynamic sharing
       : P/sec - Packets/sec, B/sec - Bytes/second, BP - Burst period (microsec)

   Rate Limiter Type   State    P/sec  P/burst      B/sec    B/burst BP      Shk
 --------------------- ----- -------- -------- ---------- ---------- ------- ---
           CEF RECEIVE   OFF        -        -          -          -       -   -
 CEF RECEIVE SECONDARY    ON    15000        -          -          - 1000000
             CEF GLEAN    ON     1000        -          -          - 1000000
            IP ERRORS   OFF        -        -          -          -       -   -
       UCAST IP OPTION    ON      100        -          -          - 1000000 G:
         ICMP ACL-DROP    ON      100        -          -          - 1000000 G:
         ICMP NO-ROUTE    ON      100        -          -          - 1000000
         ICMP REDIRECT   OFF        -        -          -          -       -   -
           RPF FAILURE    ON      100        -          -          - 1000000
          ACL VACL LOG    ON     2000        -          -          - 1000000
        ACL BRIDGED IN   OFF        -        -          -          -       -   -
       ACL BRIDGED OUT   OFF        -        -          -          -       -   -
        ARP Inspection   OFF        -        -          -          -       -   -
     DHCP Snooping IN   OFF        -        -          -          -       -   -
          IP FEATURES   OFF        -        -          -          -       -   -
           MAC PBF IN   OFF        -        -          -          -       -   -
          CAPTURE PKT   OFF        -        -          -          -       -   -
  IP ADMIS. ON L2 PORT   OFF        -        -          -          -       -   -
 MCAST IPV4 DIRECTLY C   OFF        -        -          -          -       -   -
   MCAST IPV4 FIB MISS   OFF        -        -          -          -       -   -
       MCAST IPV4 IGMP   OFF        -        -          -          -       -   -
    MCAST IPV4 OPTIONS   OFF        -        -          -          -       -   -
        MCAST IPV4 PIM   OFF        -        -          -          -       -   -
 MCAST IPV6 DIRECTLY C   OFF        -        -          -          -       -   -
        MCAST IPV6 MLD   OFF        -        -          -          -       -   -
 MCAST IPV6 CONTROL PK   OFF        -        -          -          -       -   -
          MTU FAILURE   OFF        -        -          -          -       -   -
          TTL FAILURE   OFF        -        -          -          -       -   -
 MCAST BRG FLD IP CNTR   OFF        -        -          -          -       -   -
      MCAST BRG FLD IP   OFF        -        -          -          -       -   -
             MCAST BRG   OFF        -        -          -          -       -   -
         MCAST BRG OMF   OFF        -        -          -          -       -   -
   UCAST UNKNOWN FLOOD   OFF        -        -          -          -       -   -
          LAYER_2 PDU   OFF        -        -          -          -       -   -
           LAYER_2 PT   OFF        -        -          -          -       -   -
       LAYER_2 PORTSEC   OFF        -        -          -          -       -   -
     LAYER_2 SPAN PCAP   OFF        -        -          -          -       -   -
       DIAG RESERVED 0    ON 33554431        -          -          -       1
       DIAG RESERVED 1    ON 33554431        -          -          -       1
       DIAG RESERVED 2    ON 33554431        -          -          -       1
   DIAG RESERVED LIF 0    ON 33554431        -          -          -       1
   MCAST REPL RESERVED    ON        1        -          -          -       0
Router#
```

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

```
Router# show mls rate-limit usage
P/sec - Packets/sec, B/sec - Bytes/sec, BP - Burst period (microsec), U - Usee
              Rate Limiter Type     P/sec  P/burst      B/sec    B/burst BP
            --------------------- -------- -------- ---------- ---------- ------
L3 Rate Limiters:
   RL# 1: U          ACL VACL LOG     2000        -          -          - 100000
   RL# 2: F                     -        -        -          -          -      -
   RL# 3: F                     -        -        -          -          -      -
   RL# 4: F                     -        -        -          -          -      -
```

```
   RL# 5: F                      -        -        -          -          -      -
   RL# 6: F                      -        -        -          -          -      -
   RL# 7: F                      -        -        -          -          -      -
   RL# 8: F                      -        -        -          -          -      -
   RL# 9: F                      -        -        -          -          -      -
   RL#10: U       UCAST IP OPTION        -        -      10000        100     60
                    ICMP ACL-DROP        -        -      10000        100     60
   RL#11: U          ICMP NO-ROUTE      100        -          -          - 100000
   RL#12: F                      -        -        -          -          -      -
   RL#13: F                      -        -        -          -          -      -
   RL#14: F                      -        -        -          -          -      -
   RL#15: F                      -        -        -          -          -      -
   RL#16: F                      -        -        -          -          -      -
   RL#17: F                      -        -        -          -          -      -
   RL#18: F                      -        -        -          -          -      -
   RL#19: F                      -        -        -          -          -      -
   RL#20: F                      -        -        -          -          -      -
   RL#21: F                      -        -        -          -          -      -
   RL#22: F                      -        -        -          -          -      -
   RL#23: F                      -        -        -          -          -      -
   RL#24: F                      -        -        -          -          -      -
   RL#25: F                      -        -        -          -          -      -
   RL#26: F                      -        -        -          -          -      -
   RL#27: F                      -        -        -          -          -      -
   RL#28: F                      -        -        -          -          -      -
   RL#29: F                      -        -        -          -          -      -
   RL#30: F                      -        -        -          -          -      -
   RL#31: F                      -        -        -          -          -      -

L2 Input Rate Limiters:
   RL#32: U       DIAG RESERVED 0 33554431        -          -          -      1
   RL#33: U       DIAG RESERVED 1 33554431        -          -          -      1
   RL#34: U       DIAG RESERVED 2 33554431        -          -          -      1
   RL#35: U   DIAG RESERVED LIF 0 33554431        -          -          -      1
   RL#36: U   MCAST REPL RESERVED        1        -          -          -      0
   RL#37: F                      -        -        -          -          -      -
   RL#38: F                      -        -        -          -          -      -
   RL#39: F                      -        -        -          -          -      -
   RL#40: F                      -        -        -          -          -      -
   RL#41: F                      -        -        -          -          -      -
   RL#42: F                      -        -        -          -          -      -
   RL#43: F                      -        -        -          -          -      -
   RL#44: F                      -        -        -          -          -      -
   RL#45: F                      -        -        -          -          -      -
   RL#46: F                      -        -        -          -          -      -
   RL#47: U             CEF GLEAN     1000        -          -          - 100000
   RL#48: U           RPF FAILURE      100        -          -          - 100000
   RL#49: U CEF RECEIVE SECONDARY    15000        -          -          - 100000
   RL#50: F                      -        -        -          -          -      -
   RL#51: F                      -        -        -          -          -      -

L2 Output Rate Limiters:
   RL#52: F                      -        -        -          -          -      -
   RL#53: F                      -        -        -          -          -      -
   RL#54: F                      -        -        -          -          -      -
   RL#55: F                      -        -        -          -          -      -
   RL#56: F                      -        -        -          -          -      -
   RL#57: F                      -        -        -          -          -      -
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.

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.
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#
```

Monitoring Packet Drop Statistics

Prerequisites for Packet Drop Statistics

None.

Restrictions for Packet Drop Statistics

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

Information About Packet Drop Statistics

You can use show commands to display packet drop statistics. You can capture the traffic on an interface and send a copy of this traffic to a traffic analyzer connected to a port, which can aggregate packet drop statistics.

How to Monitor Dropped Packets

Using show Commands

The PFC and DFCs 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
  MTU failures                          : 0

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

Using SPAN

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#
```

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

```
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
```

For more information, see Chapter 1, “Local SPAN, RSPAN, and ERSPAN.”

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.

For more information, see Chapter 1, "VLAN ACLs (VACLs)."

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