Configuring Denial of Service Protection

This chapter contains information on how to protect your Cisco 7600 series router against Denial of Service (DoS) attacks. The information covered in this chapter is unique to the Cisco 7600 series routers, 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 7600 Series Routers Command References*, at this URL:

  http://www.cisco.com/en/US/products/hw/routers/ps368/prod_command_reference_list.html

This chapter consists of these sections:

Understanding Control Plane Protection

To achieve routing stability, it is important to ensure that important control plane and management traffic reaches the Route Processor (RP) or Switch Processor (SP) CPU in Cisco 7600 routers. Also traffic such as IP options, Time to Live (TTL) packets, packets with errors, layer 2 broadcast and multicast packets

are forwarded to the RP and results in congestion and such traffic needs to be controlled. Protecting the control plane involves protecting the RP or SP from unnecessary traffic and also ensuring that the important control plane and management traffic is delivered to the RP or SP.

DOS protection on 7600 involves enabling the following control-plane protection mechanisms:

- Control Plane Policing (CoPP) using Modular QoS Command Line (MQC)
- Multi Layer Switching (MLS) rate limiters
- MLS protocol policing
- Broadcast and multicast storm control
- Selective Packet Discard (SPD)

By default only the basic set of control plane protection mechanisms are enabled. You should customize the control plane mechanism according to your requirement.

Implementing Control Plane Protection in Cisco 7600

This section explains the available mechanisms in the Cisco 7600 router to implement control plane protection and also explains the packet flow inside the c7600 router. Control plane protection can be implemented at different stages and it is usually applied when a packet is analyzed and forwarded to the route processor (punt packet). The decision to punt a packet is made at Policy Feature Card (PFC) or DFC level, or also at the IOS process level on RP or SP. Once the decision to punt a packet is made, you can use policing and rate limiting to implement control plane protection. Control Plane Protection can also be applied before the punt decision by configuring traffic storm control, Access Control List (ACL) or ingress Quality of Service (QoS0 in the interface and thus dropping unnecessary data packets.

The following sections explain the path of a packet towards the IOS process level on SP or RP in a Cisco 7600 router. Understanding the path of a data packet helps in configuring the control plane protection better.

Path of a Control Packet

Figure 43-1 shows the basic architecture of c7600 routers, with specifying actions in the punt path of a packet.

Figure 43-1 Path of a Punted Packet

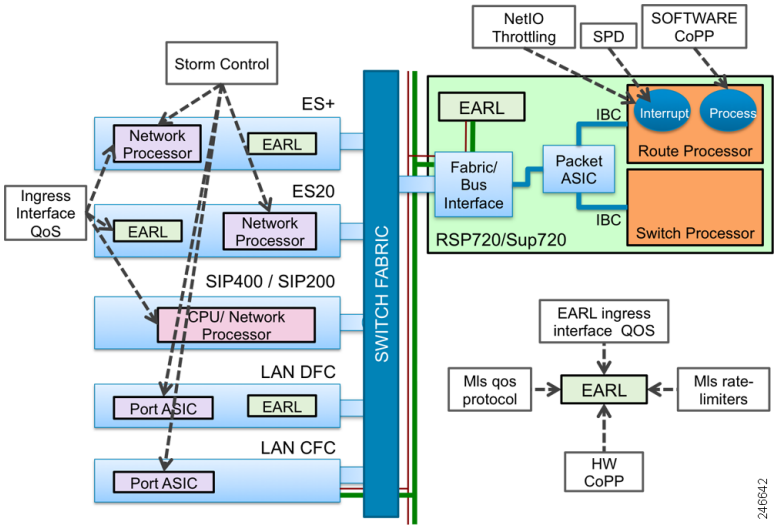

Path of a punted packet up to the forwarding decision on the PFC or DFC is similar to the path of the transit packet. When a PFC or DFC determines that the packet needs to be punted to the RP or SP CPU, the ingress line card forwards the packet to the fabric or bus interface of the active supervisor through the switch fabric. Packet then passes through the packet ASIC towards RP or SP CPU.

Linecard Support for Control Plane Protection

This section explains the QoS features and other functionalities on the supported line cards for C7600 that impacts control plane protection.

LAN Cards

Unicast, broadcast, or multicast storm-control is implemented in port Application Specific Integrated Circuits (ASICs) on LAN cards. Packets dropped as a result of storm control do not reach the EARL on DFC or PFC. Packets pass first through port ASIC and then through the Encoded Address Recognition Logic (EARL) on the LAN cards. EARL is a centralized processing engine in the c7600 supervisor engines for learning and forwarding packets based on the MAC addresses.

Ingress queuing is performed only for ports in trust Class of Service (CoS) state, and for frames that carry a valid CoS tag except for the WS-X6708 linecard where queuing can be performed on ports in DSCP state.If port is in trust CoS or DSCP state, CoS in the internal databus header is derived from the packet CoS or DSCP. Global mapping can be modified to change this setting. Otherwise, DSCP in the internal databus header is set to zero. To ensure that important packets are treated with high priority on the entire path from interface to the SP or RP CPU, it is important to preserve the original dot1q CoS or IP COS or DSCP field values in those packets. Use the **show queueing interface** command to verify the port trust state

Further actions on the packet are performed on DFC, if the line card has one, or on the PFC. Ports on SUP 720 or RSP720 can be considered as ports on a LAN card.

This example shows how to use the **show queueing interface** command to verify the port trust state.

```
Router# show queueing int Gig5/1 | inc Port
  Port QoS is enabled
```

```
Port is untrusted
```

For information on default mapping of ingress CoS to queue and Ingress DSCP queueing for WS-X6708 module, see Chapter 48, "Configuring PFC QoS".

7600-SIP-200 Cards

Unicast, broadcast or multicast storm control is not implemented on SIP-200 line cards.

QoS on the SIP-200 is configured using MQC. QoS functionalities such as classification, marking, policing, WRED, CBWFQ, shaping, priority queuing and LLQ are performed on the RX (receive) CPU on the 7600-SIP-200. QoS is not applicable for non IP packets on RX CPU.

CoS in the internal data bus header is derived from IP precedence. For non IP traffic, CoS is set to 6 for control plane traffic. For non-IP packets other than control traffic, internal CoS is set to 0. 7600-SIP-200 cards do not use the concept of port trust.

7600-SIP-400 Cards

Unicast, broadcast or multicast storm-control is not implemented on SIP-400 linecards.

When a packet is received from the SPA, Ingress Network Processor (INP) first classifies a packet. Internal **pak_priority** flag is set for IP routing protocol packets and non-IP control plane traffic (including ARP). The **pak_priority** flag in SIP-400 is has only local significance. It is considered when data bus header is built, but there is no provisioning for carrying the **pak_priority** flag to the IOS on SP or RP CPU. If **pak_priority** is set and the packet is non-IP, index-directed bit is set in the internal databus header.

The CoS field in internal databus header is copied from dot1q tag. Databus CoS value is set to 6 when packet is non-IP and **pak_priority** is set. Ingress Network Processor (INP) then applies ingress QoS if it is configured on the input interface. Packets with internal data bus CoS value set to 6 are sent towards the fabric on a separate channel. The packets are then sent to PFC through fabric interface.

7600-SIP-400 cards do not use the concept of port trust.

7600-ES20 and 7600-SIP-600 Cards

Unicast, broadcast or multicast storm control is implemented in the Network Processor (NP) on SIP-600 and ES20 linecards. Packets pass first through EARL on the SIP-600 or ES20 line card and then through the NP. On SIP-600, storm control is supported only on gigabit and tengigabit interfaces. ES20 and SIP-600 cards have a DFC. Before a packet is sent towards the EARL, databus BPDU is set for Cisco Discovery Protocol (CDP), Spanning Tree Protocol (STP), Dynamic Trunking Protocol (DTP), VLAN Trunking Protocol (VTP), Uni Directional Link Detection (UDLD), Port Aggregation Protocol (PAgP), Link Aggregation Control Protocol (LACP) or IEEE-flow-control packets. For IP packets, the databus CoS field may be copied from IP precedence, depending on the port trust state. ES20 and SIP-600 have two queues towards DFC, packets with CoS value from 0 to 5 move into the low priority queue and CoS value 6 and 7 move into the high priority queue.

7600-SIP-600 and 7600-ES20 cards use the port trust concept. Implications are the same as for LAN cards.

ES+ Line cards

Unicast, broadcast or multicast storm control is implemented in the Network Processor (NP) on the ES+ line cards. Packets pass through NP and then through the EARL on the ES+ line card. Packets dropped by storm control do not reach the EARL on the ES+ line card.

Input packet memory on Network Processor Unit (NPU) on ES+ cards has two RX queues per port (high priority and low priority queues). In addition, Intermediate System-to-Intermediate System (IS-IS) has its own special queue. High priority control packets are placed into the high priority queue.

Ingress QoS is applied in the NP. After that packets are passed to the EARL. On 7600-ES+ cards ports are always trusted.

Shared Port Adapters (SPA)

Shared Port Adapters (SPAs) are used in 7600-SIP-200, 7600-SIP-400 and 7600-SIP-600 cards.

Some SPAs may have limited QoS capabilities. For more information regarding the QoS capabilities for each SPA, see the Cisco 7600 Series Router SIP, SSC, and SPA Software Configuration Guide at:

http://www.cisco.com/en/US/docs/interfaces_modules/shared_port_adapters/configuration/7600series/sipspasw.html

Path of the Punted Packet Through Earl on PFC or DFC

This section explains the path of a punted packet through the PFC or DFC, focusing only on checks and actions that result in a punt.

When a packet is received by the PFC or DFC, lookup of the destination MAC is performed first and the appropriate action is performed depending on the type of packet as follows:

- Broadcast Frames: If there is a layer 3 interface in the VLAN on which the packet was received, broadcast frames are punted to RP CPU. Broadcast packets are not subject to MLS rate limiting.
- Multicast Frames: If there is a layer 3 interface in the VLAN on which the packet was received and if the destination MAC address corresponds to IPv4 local network control block multicast or IPv4 internetwork control block multicast (0100.5e00.0000 - 0100.5e00.01ff), the frames are punted to RP CPU. Other multicast frames may be punted to RP or SP CPU for various reasons. All packets are evaluated against the MLS QoS protocol policer and the policer is applied depending on the MLS QoS protocol policer configuration. Multicast packets punted to RP or SP CPU are subject to MLS rate limiter.
- Unicast frames: Frames with a destination MAC address of the same router may be forwarded for a layer 3 lookup. At layer 3 lookup, the result can be a punt for various reasons like layer 3 address is local to the router or feature unsupported in PFC or DFC. Unicast packets punted to the RP or SP CPU are also subject to MLS rate limiting. Unicast packets are also subject to control plane policing. You need to configure control plane policing and it is not active by default.

Policing Control Plane Traffic on the PFC or DFC

There are three mechanisms in PFC or DFC for policing control plane traffic.

- MLS protocol policer - MLS protocol policer polices all traffic of the specified protocol, not only traffic punted to the RP or SP CPU. It can also be used to change the IP precedence in these packets. It is configured using the **mls qos protocol police** command. For more information on the MLS protocol policing, see the section MLS Protocol Policing, page 43-14.

- MLS rate limiter - MLS rate limiter is applicable only to traffic punted to RP or SP CPU. EARL on PFC3 or DFC3 provides 10 rate limiter registers, which are configured in global configuration mode. These rate limiter registers contain rate limiting information for result packets that match a particular rate limiter. Out of 10 rate limiter registers, two registers are on the layer 2 forwarding engine and are reserved for layer 2 rate limiters. The remaining eight registers are in the layer 3 forwarding engine, reserved for layer 3 rate limiters. Some of the rate limiters are enabled by default and others need to be configured. MLS rate limiters are configured using the **mls rate-limit** command.

  Active and inactive rate limiters, as well as the destination index used for the punt packet can be viewed using the **show mls rate-limit** command on the RP CPU and **show mls rate-limit hw-detail** command on the SP CPU. Some rate limiting scenarios are pre-configured to share the same register. To check which rate limiters share the same register, run the **show mls rate-limit usage** command on the RP CPU. For more information on MLS rate limiter configuration, see MLS Rate-limiter Default Configuration, page 43-22
- Control plane policer - Control plane policer configured using MQC is applicable only to unicast traffic punted to RP CPU. Packet paths through the control plane policer and MLS rate limiter are mutually exclusive. If packets pass through the MLS rate limiter, they cannot be limited by the control plane policer. CoPP offers better visibility of offered rates, drop rates and packet counters. The **show policy-map control-plane** command displays the control plane policer statistics.

Figure 43-2 and Figure 43-3 describe the packet paths in EARL for unicast and multicast packets.

Figure 43-2 Packet Path in EARL for Unicast Packets

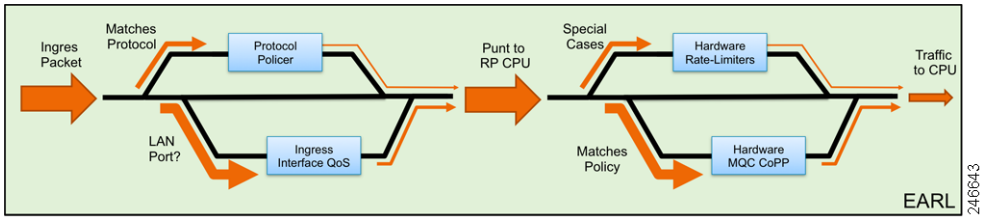

Figure 43-3 Packet Path in EARL for Multicast Packets

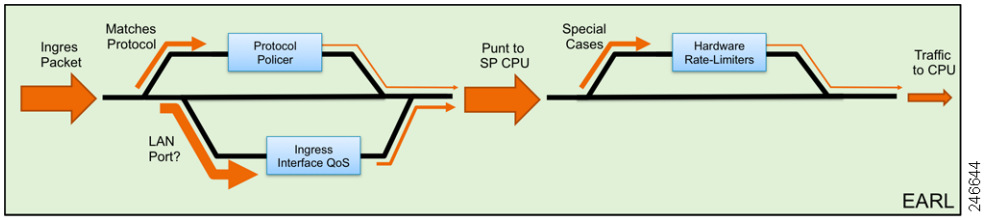

The incoming packet is first evaluated against the MLS protocol policer. MLS protocol policer and Ingress PFC or DFC QoS are mutually exclusive. MLS protocol policer and ingress PFC or DFC QoS are performed before the forwarding decision.

If forwarding decision is to punt the packet to RP CPU, the punt packet is evaluated against the configured rate limiters. If the packet matches the rate limiters, the packet passes through the appropriate rate limiter. If it does not match any of the configured rate limiters and it is destined to RP CPU, it passes through hardware CoPP. MLS rate limiters and hardware CoPP are mutually exclusive. Hardware CoPP is not applied on the punt path to SP CPU. MLS rate limiters and hardware CoPP are egress features and applied after the forwarding decision.

Figure 43-4 explains CoPP on PFC or DFC.

Figure 43-4 Control Plane Policer on PFC or DFC (Hardware CoPP)

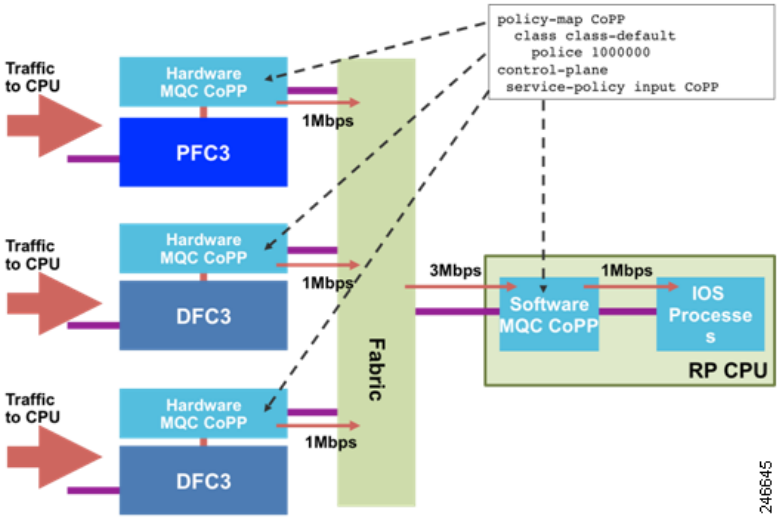

Both the MLS rate limiters and CoPP are configured in the global configuration mode. Once configured the same configuration is propagated to all PFCs or DFCs in the chassis. The first part of CoPP is applied on PFC or DFC, while the second part is applied on the RP CPU. MLS rate limiters are applied only in hardware on PFC or DFC.

Hardware policing and rate limiting for control plane protection is aggregated per PFC or DFC. The aggregate traffic coming through each PFC or DFC is again policed at the software level by CoPP.

For example, if there are two DFCs and a PFC in the chassis and aggregate policer rate in CoPP is configured as 1Mbps, the RP CPU receives a maximum of 3Mbps traffic, which is policed in the software to 1Mbps.

The command **show policy-map control-plane** command displays control plane policer statistics for all PFCs or DFCs and IOS. The **show mls rate-limit** command is executed on RP CPU shows only the configuration of the active PFC. To display the configuration of DFCs and standby PFC, execute the command **show mls rate-limit** on the DFC and standby RP respectively.

Path of the Punted Packet After PFC or DFC

RSP720/SUP 720 Hardware

Figure 43-5 describes the path of the punted packet after PFC or DFC.

Figure 43-5 Packet Path on RSP720 or SUP 720

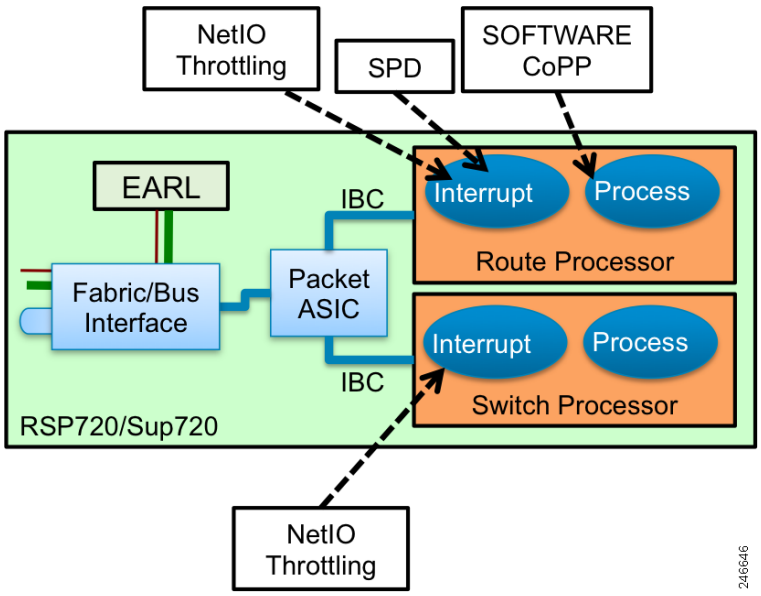

When a PFC or DFC decides to punt the packet, it instructs the ingress line card to send the packet through the switch fabric to the fabric or bus interface on the supervisor. The fabric or bus interface forwards the packet to the packet ASIC. The packet ASIC forwards the packet to the SP or RP CPU.

Each processor has a separate In-band Interface Channel (IBC) with two input queues. The high priority queue (queue 0) receives packets with the data bus CoS value from 4 to 7 and low priority queue (queue 1) receives packets with the databus CoS value from 0 to 3. When the IBC controller receives the packet, it copies the packet into IOS input/output memory and raises a Network Input/Output (NetIO) interrupt to the CPU.

For information on packets received on IBC, including the SPD statistics, use the **show ibc exec** command. This command provides information about counters, which helps to confirm that configured control plane protection is limiting the rate of punted traffic. It can be also useful in troubleshooting punted traffic on SP and RP.

This is a sample output for the **show ibc exec** command.

```
Router# show ibc | inc Rx\(
  Rx(0)   23116485        2693239358
  Rx(1)   2294534         386747297
```

Path of the Punted Packet on IOS

Once the NetIO interrupt is raised, NetIO interrupt handler routine calls the interface driver. Multiple packets received from IBC can be handled under a single NetIO interrupt. If the current NetIO interrupt is longer than the netint usec (maximum time that the network-level interrupt is allowed to run, in microseconds) period, the NetIO interrupt handler exits and NetIO interrupts are disabled during the netint mask usec (maximum time that the network-level interrupt is masked out) period. This mechanism is called NetIO throttling. During NetIO throttle, copy of packets through DMA into the IO memory happens as long as the receive buffer is not full. NetIO throttling only delays the handling of the NetIO interrupt. Since there is no control over which packets are dropped during NetIO throttling, it is important to impose proper control over packets that are punted to RP or SP CPU.

During a NetIO interrupt, the interface driver determines the type of packet and calls the the NetIO interrupt handler for the appropriate network protocol packets. If the packet is an IP packet, since software forwarding by RP CPU is permitted on c7600, IOS first calls the Cisco Express Forwarding (CEF) routine. If the packet cannot be forwarded at the interrupt level or if it is destined to the router, the packet is punted to process level.

Before CEF attempts to place the packet into the input process queue, SPD mechanism is applied. SPD allows buffering of high priority packets above the limit allowed for normal packets to avoid dropping important packets.

If the number of packets in the input queue of an interface reaches the input queue limit, interface is throttled. This means that no further packets for this interface are received until it is un-throttled. This mechanism is called interface throttling. Use the **show ibc** commands to see how many packets are dropped on throttled interfaces, and use the **show interfaces** command to see how many times the interface is throttled. To prevent drops of packets that go into SPD headroom and SPD extended headroom during interface throttle, configure the throttle selective interface configuration command.If a packet is not dropped by SPD, software CoPP is applied at the process level. Therefor it is important to make sure that CoPP is optimally configured.

Use the **show platform netint** command to check the NetIO interrupt timers and throttling statistics.

This is a sample output for the **show platform netint** command.

```
Router# show platform netint
 Network IO Interrupt Throttling:
 throttle count=90, timer count=90
 active=0, configured=1
 netint usec=4000, netint mask usec=800
inband_throttle_mask_hi = 0x0
inband_throttle_mask_lo = 0x800000
```

In the **show platform netint** command, **timer count** shows how many times **netint mask usec** timer is invoked and **throttle count** shows how many times a NetIO interrupt is throttled.

Control Plane Protection Best Practices

Control plane protection involves the following best practices:

- Deploy storm control.
- Set the port state as trust if important control plane traffic is to be received on the interface. If a port is untrusted, CoS and IP precedence are reset to zero, which affects the prioritisation of the packet in the punt path.
- Determine the CoPP policy that needs to be applied.
- Consider the limitations in different protection mechanisms on c7600 while implementing control plane protection.
- Determine if any other traffic needs to be rate limited using MLS rate limiters or MLS protocol policing.
- Monitor and adjust traffic rates. For monitoring traffic rates, these commands are useful:
  - The **show mls statistics [***module***]** command that displays information about MLS statistics.
  - The **show policy-map control-plane** command that displays information about control plane policer classification counters, especially in the class-default.
  - The **show platform netint** command that displays information regarding NetIO throttle count and timer count.
  - The **show ibc** command that displays information on classification and drop counters.
  - The **debug netdr capture** command that displays information for capturing punted packets. The captured packets can be viewed using the **show netdr captured-packets** command.

Understanding How DoS Protection Works

The following sections contain an overview of the DoS protection on the Cisco 7600 series router and describe some types of DoS attack scenarios:

- DoS Protection with a PFC3, page 43-10

DoS Protection with a PFC3

This section contains information about the available methods to counteract DoS attacks with a PFC3 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 Understanding How Control Plane Policing Works, page 43-29.

These sections describe DoS protection with a PFC3:

- Security ACLs and VACLs, page 43-11
- QoS Rate Limiting, page 43-12
- uRPF Check, page 43-12

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 router interface that is pointing to the Internet. You can apply an inbound ACL on the router 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 router interface, it matches on that ACL and drops the packet before it causes damage.

When the Cisco 7600 series router 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 on the Cisco 7600 series routers.

QoS Rate Limiting

QoS ACLs limit the amount of a particular type of traffic that is processed by the MSFC3. If a DoS attack is initiated against the MSFC, QoS ACLs can prevent the DoS traffic from reaching the MSFC 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 router from impacting the MSFC.

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 MSFC 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 FIB tables on the supervisor engine.

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 with a PFC3 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 MSFC 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 Cisco 7600 series router 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 MSFC 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 router 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 all PFC3 types.

MLS Protocol Policing

During an attack, malicious users may try to overwhelm the MSFC CPU with control packets such as routing protocol or ARP packets. These special control packets can be hardware rate limited using a specific routing protocol and an ARP policing mechanism configurable with the **mls qos protocol** command. The routing protocols supported include RIP, BGP, LDP, OSPF, IS-IS, IGRP, and EIGRP. For example, the command **mls qos protocol arp police 32000** rate limits ARP packets in hardware at 32,000 bps. Although this policing mechanism effectively protects the MSFC CPU against attacks such as line-rate ARP attacks, it does not only police routing protocols and ARP packets to the router but also polices traffic through the box 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* **pass-through** command.

This example shows how to display the available protocols to use with ARP 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 arp** command:

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

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 router. 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.

The hardware-based rate limiters available on the PFC3 are as follows:

- Ingress and egress ACL bridged packets
- uRPF check failures
- FIB receive cases
- FIB glean cases
- Layer 3 security features
- ICMP redirects
- ICMP unreachable (ACL drop)
- No-route (FIB miss)
- VACL log
- TTL failure
- MTU failure
- Multicast IPv4
- Multicast IPv6

Shared Rate-Limiters

These shared rate limiters can be configured on the Cisco 7600 router:

- IP RPF failure
- ICMP unreachable no-route
- ICMP unreachable acl-drop
- IP errors

If you enable or disbale one of the shared rate limiter, all the other shared limiters are enabled or disabled.

Recommended Rate-Limiter Configuration

The recommended rate-limiter configuration is as follows:

- Enable the rate limiters for the traffic types commonly used in a DoS attack.
- Do not use a rate limiter on VACL logging, unless you configure VACL logging.
- Disable redirects because a platform that supports hardware forwarding, such as the Cisco 7600 series router, reduces the need for redirects.
- Disable unreachables because a platform that supports hardware unreachables, such as the Cisco 7600 series router, reduces the need for 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 and so on .
  - Rate limiters do not discriminate between good frames or bad frames.
- The MTU and TTL rate limiters are enabled by default from the 15.1(1)S1 and 15.0(1)S3a release.
  - The default MTU and TTL values are 970 and 97 respectively. You can change the default values once the router is booted.
  - If non-default values on MTU and TTL have already been configured on the router, then the user defined configurations takes precedence.
  - The default values of MTU and TTL rate-limiters can be modified and saved in the configurations.
  - If the maximum supported rate limiters have already been configured, then the MTU and TTL rate limiters are not enabled by default at the boot up. The user defined rate limters is given precedence.

Ingress-Egress ACL Bridged Packets (Unicast Only)

This rate limiter rate limits packets sent to the MSFC because of an ingress/egress ACL bridge result. The router accomplishes this by altering existing and new ACL TCAM entries with a TCAM bridge result to a Layer 3 redirect result pointing to the MSFC. 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 router 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 MSFC 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 MSFC. The uRPF check rate limiters allow you to rate limit the packets per second that are bridged to the MSFC CPU when a uRPF check failure occurs.

This example shows how to rate limit the uRPF check failure packets sent to the MSFC 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 MSFC 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
```

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 MSFC). The ICMP unreachable rate limiter allows you to rate limit the packets that are sent to the MSFC containing unreachable addresses.

This example shows how to rate limit the packets that are sent to the MSFC 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 MSFC 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 MSFC. This situation occurs when traffic enters a port and contains the destination of a host on a subnet that is locally connected to the MSFC, 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 MSFC 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 MSFC 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 MSFC. For these security features, you need to rate limit the number of these packets being sent to the MSFC 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 router 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 MSFC may be overwhelmed. Rate limiting would be advantageous in this situation.

IPSec and inspection are also done by the MSFC 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 MSFC 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 router, the MSFC sends ICMP-redirect messages to the host to correct its sending path. If this traffic occurs continuously, and is not rate limited, the MSFC 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 MSFC 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 MSFC does the logging. When VACL logging is configured on the router, 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 MSFC CPU. This might cause the MSFC to be overwhelmed.

This example shows how to rate limit packets failing the MTU failures from being sent to the MSFC 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 supervisor engine. IGMP snooping listens to IGMP messages between the hosts and the supervisor engine. You cannot enable the Layer 2 PDU rate limiter if the Cisco 7600 series router is operating in truncated mode. The router uses truncated mode for traffic between fabric-enabled modules when there are both fabric-enabled and nonfabric-enabled modules installed. In this mode, the router 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 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 supervisor engine and not the MSFC CPU. You cannot enable the Layer 2 PDU rate limiter if the Cisco 7600 series router is operating in truncated mode. The router uses truncated mode for traffic between fabric-enabled modules when there are both fabric-enabled and nonfabric-enabled modules installed. In this mode, the router 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 supervisor engine. 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 Cisco 7600 series router is operating in truncated mode. The router uses truncated mode for traffic between fabric-enabled modules when there are both fabric-enabled and nonfabric-enabled modules installed. In this mode, the router 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 MSFC 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 MSFC 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 MSFC3 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 MSFC3. For this reason, it may be desirable to rate limit the flow destined to the MSFC3 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 supported with a PFC3B, PFC3BXL, PFC3C, or PFC3CXL only.

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 43-1 lists the IPv6 rate limiters and the class of traffic that each rate limiter serves.

Table 43-1 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
```

MLS Rate-limiter Default Configuration

Table 43-2 shows the MLS Rate-limiter default configuration for the PFC3 hardware-based rate limiters.

Table 43-2 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

The section contains these configuration guidelines and restrictions:

PFC3

When configuring DoS protection on systems configured with a PFC3, 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 43-29.

- Do not use these rate limiters if multicast is enabled in systems configured with a PFC3A:
  - TTL failure
  - MTU failure
- These rate limiters are supported only on a PFC3B, PFC3BXL, PFC3C, or PFC3CXL:
  - 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 ipv4 igmp/pim *rate in pps* burst *size* command to limit punting of IGMP or PIM packets in a layer 2/3 cloud.

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

The PFC3B, PFC3BXL, PFC3C, and PFC3CXL 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 router.

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.
- High CPU utilization occurs when:
  - CoPP rate limits and drops exceeding traffic
  - mls qos protocol protocol pass-through is configured

  To avoid this, rely on the CoPP to drop excessive traffic and not on **mls qos protocol** .
- 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#
```

Due to hardware limitations, PFCor DFC cannot report how many packets are passed to MLS rate limiter or dropped. However, using the **show mls statistics** and **terminal exec prompt timestamp** commands, the rate of some erroneous packets can be calculated.

MLS Rate Limiter Configuration Example

This section contains a sample MLS rate limiter configuration. To meet your requirements, you need to modify the example accordingly.

This sample configuration is valid for Cisco IOS 12.2(33)SRE3, 15.0(1)S3, 15.1(1)S2 and later releases:

```
Router(config)# mls rate-limit unicast acl input 1000 10
Router(config)# mls rate-limit unicast acl output 1000 10
Router(config)# mls rate-limit multicast ipv4 igmp 1000 10
Router(config)# mls rate-limit multicast ipv4 fib-miss 100 10
Router(config)# mls rate-limit multicast ipv4 partial 100 10
```

```
Router(config)# mls rate-limit multicast ipv4 connected 100 250
Router(config)# no mls rate-limit unicast acl vacl-log
Router(config)# mls rate-limit unicast ip options 1000 10
Router(config)# mls rate-limit multicast ipv4 ip-options 1000 10
```

For releases earlier than Cisco IOS 12.2(33)SRE3, add the following two lines to the MLS rate limiter configuration:

```
Router(config)# mls rate-limit all ttl-failure 100 10
Router(config)# mls rate-limit all mtu-failure 100 10
```

This configuration uses all the eight rate limit registers as shown in this sample output for the **show mls rate-limit** command.

```
Router# show mls rate-limit | exclude -     -
 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 DFLT ADJ   On                 100      10  Not sharing
      MCAST DIRECT CON   On                 100     250  Not sharing
        ACL BRIDGED IN   On                1000      10  Group:1 S
       ACL BRIDGED OUT   On                1000      10  Group:1 S
      MCAST PARTIAL SC   On                 100      10  Not sharing
        IP RPF FAILURE   On                 100      10  Group:0 S
           TTL FAILURE   On                  97      10  Not sharing
 ICMP UNREAC. NO-ROUTE   On                 100      10  Group:0 S
 ICMP UNREAC. ACL-DROP   On                 100      10  Group:0 S
           MTU FAILURE   On                 997      10  Not sharing
       MCAST IP OPTION   On                1000      10  Group:3 S
       UCAST IP OPTION   On                1000      10  Group:2 S
             IP ERRORS   On                 100      10  Group:0 S


Router# show mls rate-limit usage
                            Rate Limiter Type     Packets/s   Burst
                          ----------------------   ---------   -----
Layer3 Rate Limiters:
             RL# 0: Used
                                 MCAST IP OPTION        1000      10
             RL# 1: Used
                                MCAST DIRECT CON         100     250
             RL# 2: Used
                                  ACL BRIDGED IN        1000      10
                                 ACL BRIDGED OUT        1000      10
             RL# 3: Used
                                 UCAST IP OPTION        1000      10
             RL# 4: Used
                                  MCAST DFLT ADJ         100      10
             RL# 5: Used
                                  IP RPF FAILURE         100      10
                               ICMP UNREAC. NO-ROUTE         100      10
                               ICMP UNREAC. ACL-DROP         100      10
                                       IP ERRORS         100      10
             RL# 6: Used
                                     MTU FAILURE         997      10
             RL# 7: Used
                                     TTL FAILURE          97      10
             RL# 8: Rsvd for capture           -           -       -

Layer2 Rate Limiters:
             RL# 9: Reserved
```

```
RL#10: Reserved
RL#11: Free                          -           -       -
RL#12: Free                          -           -       -
```

Understanding How Control Plane Policing Works

The control plane policing (CoPP) feature increases security on the Cisco 7600 series router by protecting the MSFC from unnecessary or DoS traffic and giving priority to important control plane and management traffic. The PFC3 and DFC3 provide hardware support for CoPP. CoPP works with the PFC3 rate limiters.

The PFC3 supports the built-in "special case" rate limiters that can be used when an ACL cannot classify particular scenarios, such as IP options cases, TTL and MTU failure cases, packets with errors, and multicast packets. When enabling the special-case rate limiters, the special-case rate limiters override the CoPP policy for packets matching the rate-limiter criteria.

The traffic managed by the MSFC is divided into three functional components or *planes*:

- Data plane
- Management plane
- Control plane

The majority of traffic managed by the MSFC is handled by way of the control and management planes. You can use CoPP to protect the control and management planes, and ensure routing stability, reachability, and packet delivery. CoPP uses a dedicated control plane configuration through the modular QoS CLI (MQC) to provide filtering and rate-limiting capabilities for the control plane packets.

CoPP Default Configuration

CoPP is disabled by default and and it is recommended that you enable CoPP. For information on CoPP, see Control Plane Policing Implementation Best Practices at:
http://www.cisco.com/web/about/security/intelligence/coppwp_gs.html

CoPP Configuration Guidelines and Restrictions

When configuring CoPP, follow these guidelines and restrictions:

- Classes that match multicast are not applied in hardware but are applied in software.
- CPP is not supported in hardware for broadcast packets. The combination of ACLs, traffic storm control, and CPP software protection provides protection against broadcast DoS attacks.
- CoPP does not support ARP policies. ARP policing mechanisms provide protection against ARP storms.
- CoPP does not support non-IP classes except for the default non-IP class. ACLs can be used instead of non-IP classes to drop non-IP traffic, and the default non-IP CoPP class can be used to limit to non-IP traffic that reaches the RP CPU.
- Do not use the **log** keyword in CoPP policy ACLs.

- With PFC3A, egress QoS and CoPP cannot be configured at the same time. In this situation, CoPP is performed in the software. A warning message is displayed to inform you that egress QoS and CoPP cannot be configured at the same time.
- If you have a large QoS configuration, the system may run out of TCAM space. If this is the case, CoPP may be performed in software.
- When there is a large QoS configuration for other interfaces, you can run out of TCAM space. When this situation occurs, CoPP may be performed entirely in software and result in performance degradation and CPU cycle consumption.
- You must ensure that the CoPP policy does not filter critical traffic such as routing protocols or interactive access to the routers. Filtering this traffic could prevent remote access to the router, requiring a console connection.
- PFC3 supports built-in special-case rate limiters, which are useful for situations where an ACL cannot be used (for example, TTL, MTU, and IP options). When you enable the special-case rate limiters, you should be aware that the special-case rate limiters will override the CoPP policy for packets matching the rate-limiter criteria.
- CoPP is not enabled in hardware unless MMLS QoS is enabled globally with the **mls qos** command. If the **mls qos** command is not entered, CoPP will only work in software and will not provide any benefit to the hardware.
- Neither egress CoPP nor silent mode is supported. CoPP is only supported on ingress (service-policy output CoPP cannot be applied to the control plane interface).
- ACE hit counters in hardware are only for ACL logic. You can rely on software ACE hit counters and the **show access-list**, **show policy-map control-plane**, and **show mls ip qos** commands to troubleshoot evaluate CPU traffic.
- CoPP is performed on a per-forwarding-engine basis and software CoPP is performed on an aggregate basis.
- CoPP is not supported in hardware for multicast packets. The combination of ACLs, multicast CPU rate limiters and CoPP software protection provides protection against multicast DoS attacks.
- CoPP does not support ACEs with the **log** keyword.
- CoPP uses hardware QoS TCAM resources. Enter the **show tcam utilization** command to verify the TCAM utilization.
- When CoPP is configured and a unicast traffic passes through the CoPP classification, packets punted to RP are treated with trust DSCP action regardless of trust configured on the input port. If CoPP is configured, and you want the punted packets to be marked or trusted based on the input port, then execute the **platform ip features sequential** command on the input port. Since multicast and broadcast traffic do not go through the hardware CoPP classifcation, this is not applicable to multicast and broadcast traffic.
- When you set the policer value, note that the mls qos protocol is supported and impacts the traffic switch in the router.
- The incoming control packets needs to be trusted for prioritorizing them in control plane SPD. Otherwise, the packets may be competing with other data packets punted to RP and this increases their probability of getting dropped.
- For packets ingressing on LAN interfaces:
  - If a CoPP is not applied on the router, its preferable that either the ingress traffic DSCP or precedence is trusted using **mls qos trust**. Also avoid remarking of the incoming control protocol packets to precedence values lower than precedence 6. The control protocol packets could be classified based on their precedence or DSCP value.

  - If a CoPP is applied and a unicast traffic reaches the CoPP classification, then the CoPP overrides incoming trust state with trust dscp and preserves DSCP or precedence on the packets being punted to control plane. Multicast and broadcast traffic do not go through the hardware CoPP classification and hence this is not applicable to multicast or broadcast traffic.‘

- For packets ingressing on WAN interfaces like SIP 400 and ES+ linecards, avoid remarking the incoming control protocol packets to precedence values lower than precedence 6 or 7. The control protocol packets can be identified based on their precedence or DSCP value.
- High CPU utilization occurs when CoPP rate limits and drops the exceeding traffic and the protocol pass-through mode is configured using the **mls qos protocol** command. To avoid this, use the CoPP to drop excessive traffic and do not use the **mls qos protocol** command to route traffic directly to the route processor. For information on classifying CoPP traffic, see “Traffic Classification Overview” section on page 43-38
- CoPP processing on IPv6 traffic in the class-map also processes Layer 2 traffic.

Configuring CoPP

CoPP uses MQC to define traffic classification criteria and to specify the configurable policy actions for the classified traffic. You must first identify the traffic to be classified by defining a class map. The class map defines packets for a particular traffic class. After you have classified the traffic, you can create policy maps to enforce policy actions for the identified traffic. The **control-plane** global configuration command allows the CoPP service policies to be directly attached to the control plane.

For information on how to define the traffic classification criteria, refer to the “Defining Traffic Classification” section on page 43-38.

To configure CoPP, perform this task:

	Command	Purpose
Step 1	`Router(config)# mls qos`	Enables MLS QoS globally.
Step 2	`Router(config)# ip access-list extended` *access-list-name* `Router(config-ext-nacl)# {permit | deny}` *protocol* **source** *source-wildcard* **destination** *destination-wildcard* `[precedence` *precedence*`] [tos` *tos*`] [established] [log | log-input] [time-range` *time-range-name*`] [fragments]`	Defines ACLs to match traffic: • **permit** sets the conditions under which a packet passes a named IP access list. • **deny** sets the conditions under which a packet does not pass a named IP access list. **Note** You must configure ACLs in most cases to identify the important or unimportant traffic.
Step 3	`Router(config)# class-map` *traffic-class-name* `Router(config-cmap)# match {ip precedence} |{ip dscp} |` *access-group*	Defines the packet classification criteria. Use the **match** statements to identify the traffic associated with the class.
Step 4	`Router(config)# policy-map` *service-policy-name* `Router(config-pmap)# class` *traffic-class-name* `Router(config-pmap-c)# police {`*bits-per-second* `[`*normal-burst-bytes*`] [`*maximum-burst-bytes*`] [pir` *peak-rate-bps*`]} | [conform-action` *action*`] [exceed-action` *action*`] [violate-action` *action*`]`	Defines a service policy map. Use the **class** *traffic-class-name* command to associate classes to the service policy map. Use the **police** statements to associate actions to the service policy map.

	Command	Purpose
Step 5	`Router(config)# control-plane` `Router(config-cp)#`	Enters the control plane configuration mode.
Step 6	`Router(config-cp)# service-policy input` *service-policy-name*	Applies the QoS service policy to the control plane.

When defining the packet classification criteria, follow these guidelines and restrictions:

- To avoid matching the filtering and policing that are configured in a subsequent class, configure policing in each class. CoPP does not apply the filtering in a class that does not contain a police command. A class without a police command matches no traffic.
- The ACLs used for classification are QoS ACLs. QoS ACLs supported are IP standard, extended, and named.
- These are the only match types supported:
  - **ip precedence**
  - **ip dscp**
  - **access-group**
- Only IP ACLs are supported in hardware.
- MAC-based matching is done in software only.
- You can enter one **match** command in a single class map only.

When defining the service policy, the **police** policy-map action is the only supported action.

When applying the service policy to the control plane, the **input** direction is only supported.

Effective with Cisco IOS Release 15.3(1)S, you can configure Netflow on CoPP VLAN. Use the **ip flow ingress** command to enable the feature on control plane. The command enables all flows hitting the control plane of the router.

To configure Netflow on CoPP, perform this task:

Step 1	`Router(config)# control-plane`	Enters the control plane configuration mode.
Step 2	`Router(config-cp)# ip flow ingress`	Enables NetFlow on control plane. NetFlow will collect statistics for packets forwarded in hardware (PFC) or software (RP).
Step 3	`Router(config-cp)# no ip flow ingress`	Disables NetFlow on control plane.
Step 4	`Router(config-cp)# exit`	Exits the control plane.
Step 5	`Router(config)# exit`	Exits the configuration mode.
Step 6	`Router# show ip cache flow`	Displays the flows hitting the control-plane.

Monitoring CoPP

You can enter the **show policy-map control-plane** command for developing site-specific policies, monitoring statistics for the control plane policy, and troubleshooting CoPP. This command displays dynamic information about the actual policy applied, including rate information and the number of bytes (and packets) that conformed or exceeded the configured policies both in hardware and in software.

The output of the **show policy-map control-plane** command is as follows:

```
Router# show policy-map control-plane
Control Plane Interface
  Service policy CoPP-normal
Hardware Counters:
class-map: CoPP-normal (match-all)
      Match: access-group 130
      police :
        96000 bps 3000 limit 3000 extended limit
      Earl in slot 3 :
        0 bytes
        5 minute offered rate 0 bps
        aggregate-forwarded 0 bytes action: transmit
        exceeded 0 bytes action: drop
        aggregate-forward 0 bps exceed 0 bps
      Earl in slot 5 :
        0 bytes
        5 minute offered rate 0 bps
        aggregate-forwarded 0 bytes action: transmit
        exceeded 0 bytes action: drop
        aggregate-forward 0 bps exceed 0 bps

Software Counters:
    Class-map: CoPP-normal (match-all) 0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group 130
      police:
        96000 bps, 3125 limit, 3125 extended limit
        conformed 0 packets, 0 bytes; action: transmit
        exceeded 0 packets, 0 bytes; action: drop
        conformed 0 bps, exceed 0 bps, violate 0 bps
Router#
```

To display the hardware counters for bytes dropped and forwarded by the policy, enter the **show mls qos ip** command:

```
Router# show mls qos ip
 QoS Summary [IP]:      (* - shared aggregates, Mod - switch module)

Int Mod Dir  Class-map DSCP  Agg  Trust Fl   AgForward-By   AgPoliced-By
                             Id         Id
------------------------------------------------------------------------------
CPP  5  In CoPP-normal    0    1   dscp  0         505408       83822272
CPP  9  In CoPP-normal    0    4   dscp  0              0              0
Router#
```

To display the CoPP access list information, enter the show access-lists coppacl-bgp command:

```
Router#show access-lists coppacl-bgp
Extended IP access list coppacl-bgp
10 permit tcp host 47.1.1.1 host 10.9.9.9 eq bgp (4 matches)
20 permit tcp host 47.1.1.1 eq bgp host 10.9.9.9
30 permit tcp host 10.86.183.120 host 10.9.9.9 eq bgp (1 match)
40 permit tcp host 10.86.183.120 eq bgp host 10.9.9.9
```

Personalizng a CoPP Policy

CoPP policy applied on a router should be personalized to best fit the router traffic profile getting punted to RP.

Customizing ensures the following:

- The right kind of traffic is prioritized over other less priority or non priority traffic
- Stabilizes the network
- Achieves control plane protection
- Understand if any traffic is missed in the classification

To customize the policy, use the Mini Protocol Analyzer tool to analyze the traffic punted to RP. This tool helps you capture traffic being punted to RP and check what is the rate at which packets are punted. The data obtained can be used to identify the classes and police rates required to set up CoPP. For more information on the Mini Protocol Analyzer tool, see "Using the Mini Protocol Analyzer".

Developing a CoPP Policy

Prior to create a CoPP policy, a required volume of traffic must be identified and separated into different classes. Stratifying traffic into distinct groups based on relative importance is the recommended method:

Here are the sample classification critera used when developing CoPP policer:

- Do not use any policer in class-default .All the potential traffic should be classified in a specific class rather than in class-default.
- For catch-all traffic, use the **match ipv4 any class** or **match ipv6 any class** command. Though class-default serves the same purpose, it is recommended to minimize the traffic with class-default action as shown in this example:

```
Policy-map CoPP
   Class CLASS1
      Police <>
   Class CLASS2
     Police <>
Class MATCH-IPv4-ANY Match all IPv4 traffic which doesn't fall in any of the above
mentioned classes
     Police <>
   Class MATCH-IPV6-ANY Match all IPV6 traffic which doesn't fall in any of the above
mentioned classes
     Police <>
```

In the section "Example Of a CoPP Policy" section on page 43-37, traffic is grouped into five different classes. The actual number of classes differs and should be selected based on local requirements and security policies. These traffic classes are defined with regard to the CPU or control plane.

The five different classes are:

- Critical
  - Traffic that is crucial to the operation of the router and the network
  - Examples: routing protocols like Border Gateway Protocol (BGP)
  - Some sites might choose to classify traffic other than the ones crucial to the operation as critical when appropriate
- Important

  - Frequently used traffic that is necessary for day-to-day operations
  - Examples: traffic used for remote network access and management (telnet, Secure Shell (SSH), Network Time Protocol (NTP) and Simple Network Management Protocol (SNMP)
- Normal
  - Traffic that is functional but not essential to network operation
  - Normal traffic used to be particularly hard to address when designing control-plane protection schemes, as it should be permitted but should never pose a risk to the router. With CoPP, this traffic is permitted, but limited to a low rate.
  - Examples: ICMP echo request
- Undesirable
  - Explicitly identifies bad or malicious traffic that should be dropped and denied access to the Route Processor
  - Particularly useful when known traffic destined to the router should always be denied and not placed into a default category. Explicitly denying traffic allows the end-user to collect rough statistics on this traffic using the show commands and therefore offers some insight into the rate of denied traffic.
- Layer 2 class
  - Traffic used for address resolution protocol (ARP). Excessive ARP packets can potentially monopolize RP CPU resources, depriving other important processes of resources; CoPP can be used to rate limit ARP packets to prevent this situation. Currently, ARP is the only Layer 2 protocol specifically classified using the match protocol classification criteria.
- Match-Any class
  - Matches all the other IPv4/IPv6 traffic (which does not fall into any of the above class), and police them as appropriate. It is primarily designed so that the class-default doesn't need a policer.
- Default
  - The remaining traffic destined to the Route Processor and has not been identified
  - A default classification helps monitoring of statistics to determine the rate of unidentified traffic destined to the control-plane. The identified traffic can be further analyzed to classifyand if needed updated with the other CoPP policy entries
  - The Sup720 in Release 12.2(18)SXD1 does not support the MQC **class-default** in hardware. The support has been added effective Release 12.2(18)SXE1 software release. Anyway, this is not a big limitation as shown in the example below, where the "class-default" is replaced by a normal class-map.
  - Certain traffic types, namely Layer 2 keepalives, CLNS, and other non-IP packets will be seen by a CoPP (only in class-default). These traffic types cannot be classified by MQC for CoPP and hence, will always fall into the a CoPP class-default class. If aCoPP is configured, it is best practice to never rate limiting class-default so thatLayer 2 keepalives and other essential control plane traffic are not dropped. This is the primary reason for always configuring a "catch all" IP class in the CoPP policy-map just prior to class-default.

Using the classification scheme defined above, commonly used traffic is identified with a series of ACLs:

- Class CoPP-CRITICAL : ACL 120: critical traffic
- Class CoPP-IMPORTANT: ACL 121: important traffic

- Class CoPP-NORMAL : ACL 122: normal traffic
- Class CoPP-UNWANTED ACL 123: explicitly denies unwanted traffic (For example, slammer worm traffic)
- Class CoPP-ARP : Match the ARP protocol
- Class CoPP-Match-all ACL 124: the rest of the traffic

The ACLs build classes of traffic that are used to define the policies.

Sample CoPP Policy

This is an example of a CoPP policy developed using the ACLs above:

```
access-list 120 remark *** ACL for CoPP-Critical ***
access-list 120 remark *** LDP ***
access-list 120 permit udp 172.0.0.0 0.0.255.255 eq 646 any
access-list 120 permit udp 172.0.0.0 0.0.255.255 any eq 646
access-list 120 permit tcp 172.0.0.0 0.0.255.255 eq 646 any
access-list 120 permit tcp 172.0.0.0 0.0.255.255 any eq 646
access-list 120 remark *** BGP ***
access-list 120 permit tcp 172.0.0.0 0.0.255.255 eq bgp any
access-list 120 permit tcp 172.0.0.0 0.0.255.255 any eq bgp
access-list 120 remark *** PIM ***
access-list 120 permit pim 172.0.0.0 0.0.255.255 any
access-list 120 permit pim any 172.0.0.0 0.0.255.255
access-list 121 remark *** ACL for CoPP-IMPORTANT
access-list 121 remark *** Telnet ***
access-list 121 permit tcp 172.0.0.0 0.0.255.255 172.0.0.0 0.0.255.255 eq telnet
access-list 121 remark *** SSH ***
access-list 121 permit tcp 172.0.0.0 0.0.255.255 152.0.0.0 0.255.255.255 eq 22
access-list 121 remark *** SNMP ***
access-list 121 permit udp 172.0.0.0 0.0.255.255  152.0.0.0 0.255.255.255 eq snmp
access-list 121 remark *** NTP ***
access-list 121 permit udp 172.0.0.0 0.0.255.255  host 192.168.70.10 eq ntp
access-list 121 permit udp 172.0.0.0 0.0.255.255  host 192.168.70.30 eq ntp
access-list 121 remark *** Syslog ***
access-list 121 permit udp 172.0.0.0 0.0.255.255  152.0.0.0 0.255.255.255 eq syslog
access-list 121 remark *** TACAS+ ***
access-list 121 permit tcp 172.0.0.0 0.0.255.255  152.0.0.0 0.255.255.255 eq tacacs
access-list 122 remark CoPP normal traffic
access-list 122 permit icmp any any ttl-exceeded
access-list 122 permit icmp any any port-unreachable

access-list 122 permit icmp any any echo-reply
access-list 122 permit icmp any any echo

access-list 123 remark *** ACL for CoPP-UNDESIRABLE
access-list 123 permit icmp any any fragments
access-list 123 permit udp any any fragments
access-list 123 permit tcp any any fragments
access-list 123 permit ip any any fragments
access-list 124 remark *** ACL for CoPP-Match-all
access-list 124 permit ip any any
access-list 124 permit ipv6 any any

class-map match-all CoPP-CRITICAL
match access-group 120
class-map match-all CoPP-IMPORTANT
match access-group 121
class-map match-all CoPP-NORMAL
match access-group 122
class-map match-all CoPP-ARP
```

```
match protocol ARP
class-map match-all CoPP-UNWANTED
match access-group 123
class-map match-all CoPP-Match-all
match access-group 124
```

Although rate limiting punted traffic is recommended, ensure that the required rates of traffic are well understood, particularly for critical traffic. A very low rate might discard or drop necessary traffic, whereas a very high rate might inundate the Route Processor with non-critical packets to process. These rates are site-specific and vary depending on the local topology and routing table size.

The policed rate depends on both determined criticality and site-specific rate values. For instance, the "normal" SNMP rates differ based on environment. Using the classification scheme mentioned above, critical traffic is permitted without limitation, while important, normal, and default traffic are permitted with appropriate rate limiting. However, this deployment causes the network to drop undesirable traffic immediately.

Table 39-3 extends this example and summarizes a sample policy. Note that the rates defined in the table are used for illustrative purposes; every environment contains different baselines. For example, a large Service Provider topology would require a higher rate of critical traffic (due to large BGP routing tables) than would a typical enterprise network.

The purpose of defining the critical traffic class is not limit rates, but tag this traffic as critical and provide it with unconditional access to the Route Processor. As the policy becomes increasingly refined, a more representative rate should be used for critical traffic and show commands can detect abnormal increases in traffic rates.

Table 43-3 Sample CoPP Policy

Traffic class	Rate (bps)	Conform action	Exceed action
Critical	N/A	Transmit	Transmit
Important	125,000	Transmit	Drop
Normal	64,000	Transmit	Drop
Undesirable	32,000	Drop	Drop
ARP	64,000	Transmit	Drop
MATCH-ALL	96,000	Transmit	Drop
Class-default	-	-	-

Example Of a CoPP Policy

```
policy-map CoPP
class CoPP-CRITICAL
police 1000000 31250 31250 conform-action transmit exceed-action transmit violate-action
transmit
class CoPP-IMPORTANT
police 128000 4000 4000 conform-action transmit exceed-action drop violate-action drop
class CoPP-NORMAL
police 64000 2000 2000 conform-action transmit exceed-action drop violate-action drop
class CoPP-UNDESIRABLE
police 32000 1500 1500 conform-action transmit exceed-action drop violate-action drop
class CoPP-ARP
police 64000 1500 1500 conform-action transmit exceed-action drop violate-action drop
class CoPP-Match-all
police 96000 3000 3000 conform-action transmit exceed-action drop violate-action drop
```

```
class class-default
```

Defining Traffic Classification

The following sections contain information on how to classify CoPP traffic:

Traffic Classification Overview

You can define any number of classes, but typically traffic is grouped into classes that are based on relative importance. The following provides a sample grouping:

- Border Gateway Protocol (BGP)—Traffic that is crucial to maintaining neighbor relationships for BGP routing protocol, for example, BGP keepalives and routing updates. Maintaining BGP routing protocol is crucial to maintaining connectivity within a network or to a service provider. Sites that do not run BGP do not need to use this class.
- Interior Gateway Protocol (IGP)—Traffic that is crucial to maintaining IGP routing protocols, for example, open shortest path first OSPF, enhanced interior gateway routing protocol (EIGRP), and routing information protocol (RIP). Maintaining IGP routing protocols is crucial to maintaining connectivity within a network.
- Management—Necessary, frequently used traffic that is required during day-to-day operations. For example, traffic used for remote network access, and Cisco IOS image upgrades and management, such as telnet, secure shell (SSH), network time protocol (NTP), simple network management protocol (SNMP), terminal access controller access control system (TACACS), hypertext transfer protocol (HTTP), trivial file transfer protocol (TFTP), and file transfer protocol (FTP).
- Reporting—Traffic used for generating network performance statistics for the purpose of reporting. For example, using Cisco IOS IP service level agreements (SLAs) to generate ICMP with different DSCP settings in order to report on response times within different QoS data classes.
- Monitoring—Traffic used for monitoring a router. Traffic should be permitted but should never pose a risk to the router; with CoPP, this traffic can be permitted but limited to a low rate. For example, ICMP echo request (ping) and traceroute.
- Critical Applications—Critical application traffic that is specific and crucial to a particular customer environment. Traffic included in this class should be tailored specifically to the required application requirements of the user (in other words, one customer may use multicast, while another uses IPSec or generic routing encapsulation (GRE). For example, GRE, hot standby router protocol (HSRP), virtual router redundancy protocol (VRRP), session initiation protocol (SIP), data link switching (DLSw), dynamic host configuration protocol (DHCP), multicast source discovery protocol (MSDP), Internet group management protocol (IGMP), protocol independent multicast (PIM), multicast traffic, and IPsec.
- Layer 2 Protocols—Traffic used for address resolution protocol (ARP). Excessive ARP packets can potentially monopolize MSFC resources, starving other important processes; CoPP can be used to rate limit ARP packets to prevent this situation. Currently, ARP is the only Layer 2 protocol that can be specifically classified using the match protocol classification criteria.

- Undesirable—Explicitly identifies bad or malicious traffic that should be unconditionally dropped and denied access to the MSFC.The undesirable classification is particularly useful when known traffic destined for the router should always be denied and not placed into a default category. If you explicitly deny traffic, then you can enter **show** commands to collect approximate statistics on the denied traffic and estimate its rate.
- Default—All remaining traffic destined for the MSFC that has not been identified. MQC provides the default class, so the user can specify the treatment to be applied to traffic not explicitly identified in the other user-defined classes. This traffic has a highly reduced rate of access to the MSFC. With a default classification in place, statistics can be monitored to determine the rate of otherwise unidentified traffic destined for the control plane. After this traffic is identified, further analysis can be performed to classify it and, if needed, the other CoPP policy entries can be updated to accomodate this traffic.

After you have classified the traffic, the ACLs build the classes of traffic that are used to define the policies. For sample basic ACLs for CoPP classification, see the "Sample Basic ACLs for CoPP Traffic Classification" section on page 43-39.

Traffic Classification Guidelines

When defining traffic classification, follow these guidelines and restrictions:

- Before you develop the actual CoPP policy, you must identify and separate the required traffic into different classes. Traffic is grouped into nine classes that are based on relative importance. The actual number of classes needed might differ and should be selected based on your local requirements and security policies.
- You do not have to define policies that match bidirectionally. You only need to identify traffic unidirectionally (from the network to the MSFC) since the policy is applied on ingress only.

Sample Basic ACLs for CoPP Traffic Classification

This section shows sample basic ACLs for CoPP classification. In the samples, the commonly required traffic is identified with these ACLs:

- ACL 120—Critical traffic
- ACL 121—Important traffic
- ACL 122—Normal traffic
- ACL 123—Explicitly denies unwanted traffic
- ACL 124—All other traffic

This example shows how to define ACL 120 for critical traffic:

```
Router(config)# access-list 120 remark CoPP ACL for critical traffic
```

This example shows how to allow BGP from a known peer to this router's BGP TCP port:

```
Router(config)# access-list 120 permit tcp host 47.1.1.1 host 10.9.9.9 eq bgp
```

This example shows how to allow BGP from a peer's BGP port to this router:

```
Router(config)# access-list 120 permit tcp host 47.1.1.1 eq bgp host 10.9.9.9
Router(config)# access-list 120 permit tcp host 10.86.183.120 host 10.9.9.9 eq bgp
Router(config)# access-list 120 permit tcp host 10.86.183.120 eq bgp host 10.9.9.9
```

This example shows how to define ACL 121 for the important class:

```
Router(config)# access-list 121 remark CoPP Important traffic
```

This example shows how to permit return traffic from TACACS host:

```
Router(config)# access-list 121 permit tcp host 1.1.1.1 host 10.9.9.9 established
```

This example shows how to permit SSH access to the router from a subnet:

```
Router(config)# access-list 121 permit tcp 10.0.0.0 0.0.0.255 host 10.9.9.9 eq 22
```

This example shows how to allow full access for Telnet to the router from a host in a specific subnet and police the rest of the subnet:

```
Router(config)# access-list 121 deny tcp host 10.86.183.3 any eq telnet
Router(config)# access-list 121 permit tcp 10.86.183.0 0.0.0.255 any eq telnet
```

This example shows how to allow SNMP access from the NMS host to the router:

```
Router(config)# access-list 121 permit udp host 1.1.1.2 host 10.9.9.9 eq snmp
```

This example shows how to allow the router to receive NTP packets from a known clock source:

```
Router(config)# access-list 121 permit udp host 1.1.1.3 host 10.9.9.9 eq ntp
```

This example shows how to define ACL 122 for the normal traffic class:

```
Router(config)# access-list 122 remark CoPP normal traffic
```

This example shows how to permit router-originated traceroute traffic:

```
Router(config)# access-list 122 permit icmp any any ttl-exceeded
Router(config)# access-list 122 permit icmp any any port-unreachable
```

This example shows how to permit receipt of responses to the router that originated the pings:

```
Router(config)# access-list 122 permit icmp any any echo-reply
```

This example shows how to allow pings to the router:

```
Router(config)# access-list 122 permit icmp any any echo
```

This example shows how to define ACL 123 for the undesirable class.

```
Router(config)# access-list 123 remark explicitly defined "undesirable" traffic
```

Note In the following example, ACL 123 is a permit entry for classification and monitoring purposes, and traffic is dropped as a result of the CoPP policy.

This example shows how to permit all traffic destined to UDP 1434 for policing:

```
Router(config)# access-list 123 permit udp any any eq 1434
```

This example shows how to define ACL 124 for all other traffic:

```
Router(config)# access-list 124 remark rest of the IP traffic for CoPP
Router(config)# access-list 124 permit ip any any
```

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 router maintains ARP entries in order to forward traffic to end devices or other routers. 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 router learns the legitimate IP address with the spoofed MAC address and begins to forward traffic to that MAC address. With sticky ARP enabled, the router 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, refer to the *Cisco 7600 Series Router Cisco IOS System Message Guide* at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/core/cis7600/software/122sr/msggd/index.htm

Note Sticky ARP configurability is supported.

To configure sticky ARP on a Layer 3 interface, perform the following 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* = **ethernet**, **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#
```