



Configuring Denial of Service Protection

This chapter contains information on how to protect your system against Denial of Service (DoS) attacks. The information covered in this chapter is unique to the Catalyst 6500 series switches, and it supplements the network security information and procedures in the “[Configuring Network Security](#)” 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/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_c/index.htm
- *Cisco IOS Security Command Reference*, Release 12.2, at this URL:
http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/index.htm

This chapter consists of these sections:

- [DoS Protection Overview](#), page 24-1
- [Configuring DoS Protection](#), page 24-2

DoS Protection Overview

The DoS protection available on the Catalyst 6500 series switch provides support against two types of DoS attack scenarios:

- Data-packet processing that starves routing-protocol processing may result in DoS attacks such as the following:
 - Routing peer loss due to hello timeouts
 - HSRP peer loss due to hello timeouts
 - Routing protocol slow convergence
- Data packets congesting a CPU inband datapath may result in DoS attacks such as the following:
 - Routing peer loss due to hello packet drops
 - HSRP peer loss due to hello packet drops



Note

DoS protection used at the local router may not prevent peer loss caused by data-packet congestion on the external link.

Configuring DoS Protection

The following sections describe the different DoS protection implementations and give configuration examples:

- [Supervisor Engine DoS Protection, page 24-2](#)
- [Security ACLs, page 24-2](#)
- [QoS ACLs, page 24-4](#)
- [Forwarding Information Base Rate-Limiting, page 24-5](#)
- [ARP Throttling, page 24-5](#)
- [Monitoring Packet Drop Statistics, page 24-6](#)

Supervisor Engine DoS Protection

The supervisor engine has built-in mechanisms that limit the rate of traffic in hardware and prevent flooding of the route processor and denial of service. Rate-limiting allows most of the traffic to be dropped in hardware and only a small percentage of the traffic to be forwarded to the route processor at a nonconfigurable rate of 0.5 packets per second. Rate-limiting of packets in hardware exists for the following traffic conditions:

- ICMP unreachable messages for ACL deny

This condition allows most ACL-denied packets to be dropped in hardware, and some packets to be forwarded to the route processor for monitoring purposes.



Note Because the system is programmed to bridge all ACL-deny log packets to the route processor, we do not recommend that you configure deny log ACEs in a security ACL.

- ICMP redirect messages

ICMP redirect messages are used by routers to notify the hosts on the data link that a better route is available for a particular destination. Most of these messages are dropped in hardware and only a few messages need to reach the route processor.

- Forwarding Information Base (FIB) Failures

If the FIB does not know how to route traffic for a specific IP destination address, some packets will be forwarded to the route processor to generate ICMP redirect messages.

- Reverse Path Forwarding (RPF) Failures

If the FIB IP source address lookup results in an RPF failure, some packets will be forwarded to the route processor to generate ICMP unreachable messages.

Security ACLs

The Catalyst 6500 series switch can deny packets in hardware using security ACLs and can drop DoS packets before they reach the CPU inband datapath. Because security ACLs are applied in hardware using the TCAM, long security ACLs can be used without impacting the throughput of other traffic. Security ACLs can also be applied after a DoS attack has been identified.

When using security ACLs to drop DoS packets, note the following information:

- The security ACL must specify the traffic flow to be dropped.
- When adding a security ACL to block DoS packets to an interface that already has a security ACL configured, you must merge the DoS security ACL with the existing security ACL.
- Security ACLs need to be configured on all external interfaces that require protection. Use the interface range command to configure a security ACL on multiple interfaces.

The following example shows how a security ACL is used to drop DoS packets:

```
Router# clear mls ip mod 9
Router# show mls ip mod 9
Displaying Netflow entries in module 9
-----
DstIP          SrcIP          Prot:SrcPort:DstPort  Src i/f:AdjPtr
-----
Pkts          Bytes          Age  LastSeen  Attributes
-----
199.1.1.1      199.2.1.1      0   :0        :0        0   : 0
1843           84778          2   02:30:17  L3 - Dynamic
199.2.1.1      199.1.1.1      0   :0        :0        0   : 0
2742416       126151136     2   02:30:17  L3 - Dynamic          traffic flow identified
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# no access-list 199
Router(config)# access-list 199 deny ip host 199.1.1.1 any
Router(config)# access-list 199 permit ip any any
Router(config)# interface g9/1
Router(config-if)# ip access 199 in          security ACL applied
Router(config-if)# end
Router#
1w6d: %SYS-5-CONFIG_I: Configured from console by console
Router# clear mls ip mod 9
Router# show mls ip mod 9
Displaying Netflow entries in module 9
-----
DstIP          SrcIP          Prot:SrcPort:DstPort  Src i/f:AdjPtr
-----
Pkts          Bytes          Age  LastSeen  Attributes
-----
199.1.1.1      199.2.1.1      0   :0        :0        0   : 0
1542           70932          2   02:31:56  L3 - Dynamic
199.2.1.1      199.1.1.1      0   :0        :0        0   : 0
0              0              2   02:31:56  L3 - Dynamic          hardware-forwarded
                                                                traffic stopped

Extended IP access list 199
  deny ip host 199.1.1.1 any (100 matches)
  permit ip any any
Router# show access-list 199
Extended IP access list 199
  deny ip host 199.1.1.1 any (103 matches)          rate limiting at 0.5 pps
  permit ip any any
Router #
```

QoS ACLs

Unlike Security ACLs, QoS ACLs can be used to limit the rate of traffic without denying access to all the traffic in a flow.

When using QoS ACLs to limit the rate of packets, note the following information:

- The QoS ACL must specify the traffic flow to be rate-limited.
- When adding a QoS ACL to limit the rate of packets to an interface that already has a QoS ACL configured, you must merge the rate-limiting ACL with the existing QoS ACL.
- QoS ACLs need to be configured on all external interfaces that require protection. Use the interface range command to configure an ACL on multiple interfaces.

The following example shows how to use a QoS ACL to prevent a ping attack on a router. A QoS ACL is configured and applied on all interfaces to limit the rate of incoming ICMP echo packets.

```
Router# show ip ospf neighbors

Neighbor ID      Pri   State           Dead Time   Address        Interface
6.6.6.122        1    FULL/BDR        00:00:30   6.6.6.122     Vlan46
Router# show ip eigrp neighbors
IP-EIGRP neighbors for process 200
H   Address                Interface    Hold Uptime   SRTT   RTO  Q  Seq Type
                               (sec)          (ms)      Cnt  Num
0   4.4.4.122                V144         11 00:06:07    4     200 0 6555
Router#


ping attack starts


Router# show proc cpu | include CPU utilization
CPU utilization for five seconds: 99%/90%; one minute: 48%; five minutes: 25%
Router#
2w0d: %OSPF-5-ADJCHG: Process 100, Nbr 6.6.6.122 on Vlan46 from FULL to DOWN, Neighbor
Down: Dead timer expired
Router# show ip eigrp neighbors
IP-EIGRP neighbors for process 200
Router#
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# access-list 199 permit icmp any any echo
Router(config)# class-map match-any icmp
Router(config-cmap)# match access-group 199
Router(config-cmap)# exit
Router(config)# policy-map icmp
Router(config-pmap)# class icmp
Router(config-pmap-c)# police 96000 16000 16000 conform-action transmit exceed-action drop
Router(config-pmap-c)# exit
Router(config-pmap)# exit
Router(config)# interface range g4/1 - 9
Router(config-if-range)# service-policy input icmp
Router(config-if-range)# end


policy applied


Router(config-if-range)# end
2w0d: %SYS-5-CONFIG_I: Configured from console by console
2w0d: %OSPF-5-ADJCHG: Process 100, Nbr 6.6.6.122 on Vlan46 from LOADING to FULL, Loading
Done
Router# show ip eigrp neighbors
IP-EIGRP neighbors for process 200
H   Address                Interface    Hold Uptime   SRTT   RTO  Q  Seq Type
                               (sec)          (ms)      Cnt  Num
0   4.4.4.122                V144         13 00:00:48    8     200 0 6565
Router#
```

Forwarding Information Base Rate-Limiting

The forwarding information base (FIB) rate-limiting allows all packets that require software processing to be rate limited.

The following FIB rate-limiting usage guidelines apply:

- FIB rate-limiting does not limit the rate of multicast traffic.
- FIB rate-limiting does not differentiate between legitimate and illegitimate traffic (for example, tunnels, Telnet).
- FIB rate-limiting applies aggregate rate-limiting and not per flow rate-limiting.

The following example shows traffic destined for a nonexistent host address on a locally connected subnet. Normally, the ARP request would result in an ARP reply and the installation of a FIB adjacency for this traffic. However, the adjacency in the FIB for the destination subnet would continue to receive traffic that would, in turn, be forwarded for software processing. By applying rate-limiting to this traffic, the rate of traffic forwarded for software processing can be limited to a manageable amount.

```
Router# show ip eigrp neighbors
IP-EIGRP neighbors for process 200
H   Address                Interface    Hold Uptime    SRTT   RTO   Q   Seq Type
                               (sec)          (ms)          Cnt Num
0   4.4.4.122                V144        11 00:00:26    8     200   0   6534
Router# show ip ospf neighbors

Neighbor ID      Pri   State           Dead Time   Address        Interface
6.6.6.122        1    FULL/BDR        00:00:36   6.6.6.122     Vlan46
→ Router#                               attack starts
Router# show arp | include 199.2.250.250
Internet 199.2.250.250      0    Incomplete     ARPA
Router#
1w6d: %OSPF-5-ADJCHG: Process 100, Nbr 6.6.6.122 on Vlan46 from FULL to DOWN, Neighbor
Down: Dead timer expired
Router# show ip eigrp neighbors
IP-EIGRP neighbors for process 200
Router#
Router# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
→ Router(config)# mls ip cef rate-limit 1000          traffic rate limited to 1000 pps
Router(config)# end
Router#
1w6d: %SYS-5-CONFIG_I: Configured from console by console
Router#
1w6d: %OSPF-5-ADJCHG: Process 100, Nbr 6.6.6.122 on Vlan46 from LOADING to FULL, Loading
Done
Router# show ip eigrp neighbors
IP-EIGRP neighbors for process 200
H   Address                Interface    Hold Uptime    SRTT   RTO   Q   Seq Type
                               (sec)          (ms)          Cnt Num
0   4.4.4.122                V144        12 00:00:07    12    200   0   6536
Router#
```

ARP Throttling

ARP throttling limits the rate at which packets destined to a connected network are forwarded to the route processor. Most of these packets are dropped, but a small number are sent to the router (rate limited).

Monitoring Packet Drop Statistics

Because the rate-limiting mechanism allows a certain number of packets to be forwarded for software processing, you can view the packet drop statistics by entering NetFlow **show** commands from the CLI. You can also capture the incoming or outgoing traffic on an interface and send a copy of this traffic to an external interface for monitoring by, for example, a traffic analyzer. To capture traffic and forward it to an external interface, use the **monitor session** commands.

Monitoring Dropped Packets Using NetFlow Commands

The following NetFlow commands display flows that are destined to the router MAC that are either hardware switched or forwarded to the route processor.

Displaying statistics based on source or flow only works if the MLS NetFlow flowmask is set to a value greater than destination-only.

```
Router# show mls ip
Displaying Netflow entries in Supervisor Earl
-----
DstIP          SrcIP          Prot:SrcPort:DstPort  Src i/f:AdjPtr
-----
200.2.5.3      0.0.0.0        0 :0 :0          0 : 0

Pkts          Bytes          Age  LastSeen  Attributes
-----
0             0              1   01:52:25  L3 - Dynamic
```

```
Router# show mls netflow flowmask
current ip flowmask for unicast: destination only
current ipx flowmask for unicast: destination only
```

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# mls flow ip destination-source
Router(config)# exit
1w6d: %SYS-5-CONFIG_I: Configured from console by console
Router# show mls ip
Displaying Netflow entries in Supervisor Earl
-----
DstIP          SrcIP          Prot:SrcPort:DstPort  Src i/f:AdjPtr
-----
200.2.5.3      223.255.254.226 0 :0 :0          0 : 0

Pkts          Bytes          Age  LastSeen  Attributes
-----
0             0              2   01:54:05  L3 - Dynamic
```

```
Router#
```

When you use the **show mls ip** command to display information about flows for a specific source or destination address, the command accepts 32 host prefixes only. When you use the output modifiers, you might see all flows from a specific subnet.

```
Router# show mls ip source 9.9.9.2 mod 4
Displaying Netflow entries in module 4
-----
DstIP          SrcIP          Prot:SrcPort:DstPort  Src i/f:AdjPtr
-----
9.9.9.177      9.9.9.2        0 :0 :0          0 : 0

Pkts          Bytes          Age  LastSeen  Attributes
-----
0             0              28  01:56:59  L3 - Dynamic

Router# show mls ip mod 4 | include 9.9.9
9.9.9.177      9.9.9.2        0 :0 :0          0 : 0
```

```
9.9.9.177      9.9.9.1      0   :0      :0      0   : 0
```

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
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
```

