Configuring IPv4 and IPv6 Access Control Lists

Cisco MDS 9000 Series Switches can route IP version 4 (IPv4) traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature routes traffic between VSANs. To do so, each VSAN must be in a different IPv4 subnetwork. Each Cisco MDS 9000 Series Switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

IPv4 Access Control Lists (IPv4-ACLs and IPv6-ACLs) provide basic network security to all Cisco MDS 9000 Series Switches. IPv4-ACLs and IPv6-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each Cisco MDS 9000 Series Switch can have a maximum total of 128 IPv4-ACLs or 128 IPv6-ACLs and each IPv4-ACL or IPv6-ACL can have a maximum of 256 filters.

This chapter includes the following sections:

- About IPv4 and IPv6 Access Control Lists, on page 2
- IPv4-ACL and IPv6-ACL Configuration Guidelines, on page 2
- About Filter Contents, on page 3
- Creating IPv4-ACLs or IPv6-ACLs, on page 6
- Creating IPv4-ACLs, on page 6
- Creating IPv6-ACLs, on page 7
- Defining IPv4-ACLs, on page 7
- Defining IPv6-ACLs, on page 8
- Operand and port options for an IPv4-ACL, on page 8
- Operand and port options for an IPv6-ACL, on page 9
- Adding IP Filters to an Existing IPv4-ACL, on page 9

- Adding IP Filters to an Existing IPv6-ACL, on page 10
- Removing IP Filters from an Existing IPv4-ACL, on page 10
- Removing IP Filters from an Existing IPv6-ACL, on page 11
- Verifying the IPv4-ACL or IPv6-ACL Configuration, on page 11
- Reading the IP-ACL Log Dump, on page 12
- Applying an IP-ACL to an Interface, on page 13
- Applying an IPv6-ACL to an Interface, on page 15
- Applying an IP-ACL to mgmt0, on page 15
- Open IP Ports on Cisco MDS 9000 Series Platforms, on page 16
- IP-ACL Counter Cleanup, on page 18

About IPv4 and IPv6 Access Control Lists

Cisco MDS 9000 Family switches can route IP version 4 (IPv4) traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature routes traffic between VSANs. To do so, each VSAN must be in a different IPv4 subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

IPv4 Access Control Lists (IPv4-ACLs and IPv6-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4-ACLs and IPv6-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum total of 128 IPv4-ACLs or 128 IPv6-ACLs and each IPv4-ACL or IPv6-ACL can have a maximum of 256 filters.

IPv4-ACL and IPv6-ACL Configuration Guidelines

Follow these guidelines when configuring IPv4-ACLs or IPv6-ACLs in any switch or director in the Cisco MDS 9000 Family:

- You can apply IPv4-ACLs or IPv6-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Caution If IPv4-ACLs or IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group. Do not apply IPv4-ACLs or IPv6-ACLs to only one member of a PortChannel group. Apply IPv4-ACLs or IPv6-ACLs to the entire channel group.

- Configure the order of conditions accurately. As the IPv4-ACL or the IPv6-ACL filters are sequentially applied to the IP flows, only the first match determines the action taken. Subsequent matches are not considered. Be sure to configure the most important condition first. If no conditions match, the software drops the packet.
- Configure explicit deny on the IP Storage Gigabit Ethernet ports to apply IP ACLs because implicit deny does not take effect on these ports.

About Filter Contents

An IP filter contains rules for matching an IP packet based on the protocol, address, port, ICMP type, and type of service (TS).

This section includes the following topics:

Protocol Information

The protocol information is required in each filter. It identifies the name or number of an IP protocol. You can specify the IP protocol in one of two ways:

- Specify an integer ranging from 0 to 255. This number represents the IP protocol.
- Specify the name of a protocol including, but not restricted to, Internet Protocol (IP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP).

Note When configuring IPv4-ACLs or IPv6-ACLs on Gigabit Ethernet interfaces, only use the TCP or ICMP options.

Address Information

The address information is required in each filter. It identifies the following details:

- Source—The address of the network or host from which the packet is being sent.
- Source-wildcard—The wildcard bits applied to the source.
- Destination—The number of the network or host to which the packet is being sent.
- Destination-wildcard—The wildcard bits applied to the destination.

Specify the source and source-wildcard or the destination and destination-wildcard in one of two ways:

- Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
  - Each wildcard bit set to zero indicates that the corresponding bit position in the packet's IPv4 address must exactly match the bit value in the corresponding bit position in the source.
  - Each wildcard bit set to one indicates that both a zero bit and a one bit in the corresponding position of the packet's IPv4 or IPv6 address will be considered a match to this access list entry. Place ones in the bit positions you want to ignore. For example, 0.0.255.255 requires an exact match of only the first 16 bits of the source. Wildcard bits set to one do not need to be contiguous in the source-wildcard. For example, a source-wildcard of 0.255.0.64 would be valid.

- Using the **any** option as an abbreviation for a source and source-wildcard or destination and destination-wildcard (0.0.0.0/255.255.255.255)

Port Information

The port information is optional. To compare the source and destination ports, use the **eq** (equal) option, the **gt** (greater than) option, the **lt** (less than) option, or the **range** (range of ports) option. You can specify the port information in one of two ways:

- Specify the number of the port. Port numbers range from 0 to 65535. The following table displays the port numbers recognized by the Cisco NX-OS software for associated TCP and UDP ports.
- Specify the name of a TCP or UDP port as follows:
  - TCP port names can only be used when filtering TCP.
  - UDP port names can only be used when filtering UDP.

Table 1: TCP and UDP Port Numbers

Protocol	Port	Number
UDP	dns	53
	dhcps	67
	tftp	69
	rpcbind	111
	ntp	123
	radius accounting	1646 or 1813
	radius authentication	1645 or 1812
	snmp	161
	snmp-trap	162
	syslog	514
	nfs	2049

Protocol	Port	Number
TCP[1]	ftp	20
	ftp-data	21
	ssh	22
	telnet	23
	smtp	25
	tasacs-ds	65
	www	80
	sftp	115
	http	143
	ldap no secure	389
	https	443
	ldap secure	636
	wbem-http	5988
	wbem-https	5989

[1] If the TCP connection is already established, use the established option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set.

ICMP Information

IP packets can be filtered based on the following optional ICMP conditions:

- icmp-type—The ICMP message type is a number from 0 to 255.
- icmp-code—The ICMP message code is a number from 0 to 255.

The following table displays the value for each ICMP type.

Table 2: ICMP Type Value

ICMP Type [2]	Code
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

[2] ICMP redirect packets are always rejected.

ToS Information

IP packets can be filtered based on the following optional ToS conditions:

- ToS level—The level is specified by a number from 0 to 15.
- ToS name—The name can be max-reliability, max-throughput, min-delay, min-monetary-cost, and normal.

Creating IPv4-ACLs or IPv6-ACLs

Traffic coming into the switch is compared to IPv4-ACL or IPv6-ACL filters based on the order that the filters occur in the switch. New filters are added to the end of the IPv4-ACL or the IPv6-ACL. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the filter, the traffic is denied. For this reason, you should have the frequently hit filters at the top of the filter. There is an *implied deny* for traffic that is not permitted. A single-entry IPv4-ACL or IPv6-ACL with only one deny entry has the effect of denying all traffic.

To configure an IPv4-ACL or an IPv6-ACL, follow these steps:

Procedure

Step 1 Create an IPv4-ACL or an IPv6-ACL by specifying a filter name and one or more access condition(s). Filters require the source and destination address to match a condition. Use optional keywords to configure finer granularity.

Note The filter entries are executed in sequential order. You can only add the entries to the end of the list. Take care to add the entries in the correct order.

Step 2 Apply the access filter to specified interfaces.

Creating IPv4-ACLs

To create an IPv4-ACL, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ip access-list List1 permit ip any any**

Configures an IPv4-ACL called List1 and permits IP traffic from any source address to any destination address.

Step 3 switch(config)# **no ip access-list List1 permit ip any any**

(Optional) Removes the IPv4-ACL called List1.

Step 4 switch(config)# **ip access-list List1 deny tcp any any**

Updates List1 to deny TCP traffic from any source address to any destination address.

Creating IPv6-ACLs

To create an IPv6-ACL, follow these steps:

Procedure

Step 1 switch# **configure terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ipv6 access-list List1**

switch(config-ipv6-acl)#

Configures an IPv6-ACL called List1 and enters IPv6-ACL configuration submode.

Step 3 switch(config)# **no ipv6 access-list List1**

(Optional) Removes the IPv6-ACL called List1 and all its entries.

Step 4 switch(config-ipv6-acl)# **permit ipv6 any any**

Adds an entry permitting IPv6 traffic from any source address to any destination address.

Step 5 switch(config-ipv6-acl)# **no permit ipv6 any any**

(Optional) Removes an entry from the IPv6-ACL.

Step 6 switch(config-ipv6-acl)# **deny tcp any any**

Adds an entry to deny TCP traffic from any source address to any destination address.

Defining IPv4-ACLs

To define an IPv4-ACL that restricts management access, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ip access-list restrict_mgmt permit ip 10.67.16.0 0.0.0.255 any**

Defines an entry in an IPv4-ACL named restrict_mgmt allowing all addresses in the 10.67.16.0/24 subnet.

Step 3 switch(config)# **ip access-list restrict_mgmt permit icmp any any eq 8**

Adds an entry to an IPv4-ACL named restrict_mgmt to allow any device to ping the MDS (icmp type 8).

Step 4 switch(config)# **ip access-list restrict_mgmt deny ip any any**

Explicitly blocks all other access to an access-list named restrict_mgmt.

Defining IPv6-ACLs

To define an IPv6-ACL that restricts management access, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ip access-list RestrictMgmt**

switch(config-ipv6-acl)#

Configures an IPv6-ACL called RestrictMgmt and enters IPv6-ACL configuration submode.

Step 3 switch(config)# **permit ipv6 2001:0DB8:800:200C::/64 any**

Defines an entry allowing all addresses in the 2001:0DB8:800:200C::/64 prefix.

Step 4 switch(config)# **permit icmp any any eq 8**

Adds an entry to allow any device to ping the MDS (ICMP type 8).

Step 5 switch(config)# **deny ipv6 any any**

Explicitly blocks all other IPv6 access.

Operand and port options for an IPv4-ACL

To use the operand and port options for an IPv4-ACL, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any**

Denies TCP traffic from 1.2.3.0 through source port 5 to any destination.

Operand and port options for an IPv6-ACL

To use the operand and port options for an IPv6-ACL, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ip access-list List2 deny tcp 2001:0DB8:800:200C::/64 eq port 5 any**

Denies TCP traffic from 2001:0DB8:800:200C::/64 through source port 5 to any destination.

Adding IP Filters to an Existing IPv4-ACL

After you create an IPv4-ACL or an IPv6-ACL, you can add subsequent IP filters at the end of the IPv4-ACL or the IPv6-ACL. You cannot insert filters in the middle of an IPv4-ACL or an IPv6-ACL. Each configured entry is automatically added to the end of a IPv4-ACL or a IPv6-ACL.

To add entries to an existing IPv4-ACL, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **ip access-list List1 permit tcp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0 eq port telnet**

Permits TCP for Telnet traffic.

Step 3 switch(config)# **ip access-list List1 permit tcp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0 eq port http**

Permits TCP for HTTP traffic.

Step 4 switch(config)# **ip access-list List1 permit udp 10.1.1.2 0.0.0.0 172.16.1.1 0.0.0.0**

Permits UDP for all traffic.

Adding IP Filters to an Existing IPv6-ACL

To add entries to an existing IPv6-ACL, follow these steps:

Procedure

Step 1 switch# **configure terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ipv6 access-list List2**

switch(config-ipv6-acl)#

Configures an IPv6-ACL and enters IPv6-ACL configuration submode.

Step 3 switch(config-ipv6-acl)# **permit ip 2001:0DB8:800:200C::/64 2001:0DB8:800:2010::/64 eq 23**

Permits TCP for Telnet traffic.

Step 4 switch(config-ipv6-acl)# **permit tcp 2001:0DB8:800:200C::/64 2001:0DB8:800:2010::/64 eq 143**

Permits TCP for HTTP traffic.

Step 5 switch(config-ipv6-acl)# **permit udp 2001:0DB8:800:200C::/64 2001:0DB8:800:2010::/64**

Permits UDP for all traffic.

Removing IP Filters from an Existing IPv4-ACL

To remove configured entries from an IPv4-ACL, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no ip access-list List2 deny tcp 1.2.3.0 0.0.0.255 eq port 5 any**

Removes this entry from the IPv4-ACL (List2).

Step 3 switch(config)# **no ip access-list x3 deny ip any any**

Removes this entry from the IPv4-ACL (x3).

Step 4 switch(config)# **no ip access-list x3 permit ip any any**

Removes this entry from the IPv4-ACL (x3).

Removing IP Filters from an Existing IPv6-ACL

To remove configured entries from an IPv6-ACL, follow these steps:

Procedure

Step 1 switch# **configure terminal**

switch(config)#

Enters configuration mode.

Step 2 switch(config)# **ipv6 access-list List3**

switch(config-ipv6-acl)#

Configures an IPv6-ACL and enters IPv6-ACL configuration submode.

Step 3 switch(config-ipv6-acl)# **no deny tcp 2001:0DB8:800:2010::/64 eq port 5 any**

Removes the TCP entry from the IPv6-ACL.

Step 4 switch(config-ipv6-acl)# **no deny ip any any**

Removes the IP entry from the IPv6-ACL.

Verifying the IPv4-ACL or IPv6-ACL Configuration

Use the **show ip access-list** command to view the contents of configured IPv4-ACLs. An IPv4-ACL can have one or more filters. (See the following examples).

Displays Filters Configured for an IPv4-ACL

```
switch# show ip access-list abc

ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (2 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (7 matches)
```

Displays Configured IPv6-ACLs

Use the **show ipv6 access-list** command to view the contents of configured access filters. Each access filter can have several conditions. (See the following examples).

```
switch# show ipv6 access-list
```

```
switch# show ipv6 access-list
IPv6 access list copp-system-acl-bgp6
        10 permit tcp any gt 1024 any eq bgp
        20 permit tcp any eq bgp any gt 1024
IPv6 access list copp-system-acl-icmp6
        10 permit icmp any any echo-request
        20 permit icmp any any echo-reply
IPv6 access list copp-system-acl-icmp6-msgs
        10 permit icmp any any router-advertisement
        20 permit icmp any any router-solicitation
        30 permit icmp any any nd-na
        40 permit icmp any any nd-ns
        50 permit icmp any any mld-query
        60 permit icmp any any mld-report
        70 permit icmp any any mld-reduction
IPv6 access list copp-system-acl-ntp6
        10 permit udp any any eq ntp
        20 permit udp any eq ntp any
IPv6 access list copp-system-acl-ospf6
        10 permit 89 any any
IPv6 access list copp-system-acl-pim6
        10 permit 103 any ff02::d/128
        20 permit udp any any eq pim-auto-rp
IPv6 access list copp-system-acl-radius6
```

Displays a Summary of the Specified IPv6-ACL

```
switch# show ipv6 access-list abc
```

Reading the IP-ACL Log Dump

Use the LogEnabled check box option during IP filter creation to log information about packets that match this filter. The log output displays the ACL number, permit or deny status, and port information.

Use the **log-deny** option at the end of a filter condition to log information about packets that match dropped entries. The log output displays the ACL number, permit or deny status, and port information.

Note To capture these messages in a logging destination, you must configure severity level 7 for the kernel and ipacl facilities and severity level 7 for the logging destination: logfile, monitor.

```
switch# configure terminal
switch(config)# logging level kernel 7
switch(config)# logging level ipacl 7
switch(config)# logging logfile message 7
```

For the input ACL, the log displays the raw MAC information. The keyword “MAC=” does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not logged.

The following example is an input ACL log dump:

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:00:08:00
:45:00:00:54:00:00:40:00:40:01:0e:86:0b:0b:0b:0c:0b:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02
:01:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b
:1c:1d:1e:1f:20:21:22:23:24:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84
TOS=0x00
PREC=0x00 TTL=64 ID=0 DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

The following example is an output ACL log dump:

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2 DST=11.11.11.12 LEN=84 TOS=0x00 PREC=0x00 TTL=255
 ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277 SEQ=1280
```

Applying an IP-ACL to an Interface

You can define IP-ACLs without applying them. However, the IP-ACLs will have no effect until they are applied to an interface on the switch. You can apply IP-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Tip Apply the IP-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IPv4-ACL to M0 on Switch 1 instead of an outbound filter to M1 on Switch 3 (See Figure 1: Denying Traffic on the Inbound Interface, on page 13.

Figure 1: Denying Traffic on the Inbound Interface

The **access-group** option controls access to an interface. Each interface can only be associated with one IP-ACL per direction. The ingress direction can have a different IP-ACL than the egress direction. The IP-ACL becomes active when applied to the interface.

Tip Create all conditions in an IP-ACL before applying it to the interface.

Caution If you apply an IP-ACL to an interface before creating it, all packets in that interface are dropped because the IP-ACL is empty.

The terms *in*,*out*,*source* , and *destination* are used as referenced by the switch:

- In—Traffic that arrives at the interface and goes through the switch; the source is where it transmitted from and the destination is where it is transmitted to (on the other side of the router).

Tip The IP-ACL applied to the interface for the ingress traffic affects both local and remote traffic.

- Out—Traffic that has already been through the switch and is leaving the interface; the source is where it transmitted from and the destination is where it is transmitted to.

Tip The IP-ACL applied to the interface for the egress traffic only affects local traffic.

To apply an IPv4-ACL to an interface, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface mgmt0**

switch(config-if)#

Configures a management interface (mgmt0).

Step 3 switch(config-if)# **ip access-group restrict_mgmt**

Applies an IPv4-ACL called restrict_mgmt for both the ingress and egress traffic (default).

Step 4 switch(config-if)# **no ip access-group NotRequired**

Removes the IPv4-ACL called NotRequired.

Step 5 switch(config-if)# **ip access-group restrict_mgmt in**

Applies an IPv4-ACL called restrict_mgmt (if it does not already exist) for ingress traffic.

Step 6 switch(config-if)# **no ip access-group restrict_mgmt in**

Removes the IPv4-ACL called restrict_mgmt for ingress traffic.

Step 7 switch(config-if)# **ip access-group SampleName2 out**

Applies an IPv4-ACL called SampleName2 (if it does not already exist) for local egress traffic.

Step 8 switch(config-if)# **no ip access-group SampleName2 out**

Removes the IPv4-ACL called SampleName2 for egress traffic.

Applying an IPv6-ACL to an Interface

To apply an IPv6-ACL to an interface, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface mgmt0**

switch(config-if)#

Configures a management interface (mgmt0).

Step 3 switch(config-if)# **ipv6 traffic-filter RestrictMgmt in**

Applies an IPv6-ACL called RestrictMgmt (if it does not already exist) for ingress traffic.

Step 4 switch(config-if)# **no ipv6 traffic-filter RestrictMgmt in**

Removes the IPv6-ACL called RestrictMgmt for ingress traffic.

Step 5 switch(config-if)# **ipv6 traffic-filter SampleName2 out**

Applies an IPv6-ACL called SampleName2 (if it does not already exist) for egress traffic.

Step 6 switch(config-if)# **no ipv6 traffic-filter SampleName2 out**

Removes the IPv6-ACL called SampleName2 for egress traffic.

Applying an IP-ACL to mgmt0

A system default ACL called mgmt0 exists on the mgmt0 interface. This ACL is not visible to the user, so mgmt0 is a reserved ACL name that cannot be used. The mgmt0 ACL blocks most ports and only allows access to required ports in compliance to accepted security policies.

Note If you apply an ACL to the mgmt0 interface, it automatically replaces the system default ACL on the mgmt0 interface. When you remove the user-defined ACL on the mgmt0 interface, system automatically reapplies the mgmt0 to the system default ACL. We recommend that you configure an ACL to open only the ports that are required and deny the ports that are not required.

Verifying Interface IP-ACL Configuration

Use the **show interface** command to display the IPv4-ACL configuration on an interface.

```
switch# show interface mgmt 0
mgmt0 is up
    Internet address(es):
        10.126.95.180/24
        2001:420:54ff:a4::222:5dd/119
        fe80::eaed:f3ff:fee5:d28f/64
    Hardware is GigabitEthernet
    Address is e8ed.f3e5.d28f
    MTU 1500 bytes, BW 1000 Mbps full Duplex
    5144246 packets input, 1008534481 bytes
      2471254 multicast frames, 0 compressed
      0 input errors, 0 frame
      0 overrun, 0 fifo
    1765722 packets output, 1571361034 bytes
      0 underruns, 0 output errors
      0 collisions, 0 fifo
      0 carrier errors
```

Use the **show interface** command to display the IPv6-ACL configuration on an interface.

```
switch# show interface gigabitethernet 2/1

GigabitEthernet2/1 is up
Hardware is GigabitEthernet, address is 000e.38c6.28b0
Internet address is 10.1.1.10/24
MTU 1500 bytes
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
Auto-Negotiation is turned on
ip access-group RestrictMgmt
5 minutes input rate 1208 bits/sec, 151 bytes/sec, 2 frames/sec
5 minutes output rate 80 bits/sec, 10 bytes/sec, 0 frames/sec
6232 packets input, 400990 bytes
0 multicast frames, 0 compressed
0 input errors, 0 frame, 0 overrun 0 fifo
503 packets output, 27054 bytes, 0 underruns
0 output errors, 0 collisions, 0 fifo
0 carrier errors
```

Open IP Ports on Cisco MDS 9000 Series Platforms

Cisco MDS 9000 Series platforms with default configurations have IP ports that are open on the external management interface. The table below lists the open ports and their corresponding services:

Table 3: Open IP Ports on Cisco MDS 9000 Series Platforms

Port number	IP Protocol (UDP/TCP)	Platform	Feature/Service Name	Random Port?
None	UDP	All	—	—
600 - 1024	TCP	All	NFS	Yes
2002	TCP	All	Remote Packet Capture	No
7546	TCP	All	CFS over IPv4	No

Port number	IP Protocol (UDP/TCP)	Platform	Feature/Service Name	Random Port?
9333	TCP	All	Cluster	No
32768 - 32769	TCP	Cisco MDS 8-Gb Fabric Switch for HP c-Class Blade System Cisco MDS 9148 Cisco MDS 9222i Cisco MDS 9506 Cisco MDS 9509 Cisco MDS 9513	License Manager	Yes
44583 - 59121	TCP	Cisco MDS 9148S Cisco MDS 9250i Cisco MDS 9706 Cisco MDS 9710	License Manager	Yes

NFS—A port in this range is used by the NFS service on the switch. This is only for intraswitch use. It is not essential to provide external access to or from these ports. This feature cannot be disabled. To block access to this service, configure an IP access list to deny access to the range of ports. Refer to the About IPv4 and IPv6 Access Control Lists section for more details.

Remote Packet Capture—This port is used by the Fibre Channel Analyzer service on the switch for communicating with an Ethereal protocol analyzer client on a host using the Remote Capture Protocol (RPCAP). This service is used for troubleshooting and is optional for normal switch operation. This feature cannot be disabled. To block access to this service, configure an IP access list to deny access to the range of ports. Refer to the About IPv4 and IPv6 Access Control Lists section for more details.

CFS over IPv4—This port is used by the CFS over IPv4 service to distribute switch configuration information to peer switches in the fabric. CFS is an important service for a switch to communicate with peers, but several transport options are possible. The correct transport depends on the fabric implementation. This port may be closed by disabling the CFS over IPv4 service. Refer to the Enabling CFS Over IP section of the *Cisco MDS 9000 Family CLI Configuration Guide* for details.

Cluster—This port is used by the cluster service to communicate with peer switches in a cluster. Features such as IOA and SME rely on this service. If such features are not in use, the cluster service is not essential to a switch operation. This port can be closed by disabling the cluster service. Refer to the Enabling and Disabling Clustering section of the *Cisco MDS 9000 Family Storage Media Encryption Configuration Guide* for details.

License Manager—These ports are used by the License Manager service. This only for intraswitch use. It is not essential to provide external access to or from these ports. This feature cannot be disabled. To block access to this service, configure an IP access list to deny access to the range of ports. Refer to the About IPv4 and IPv6 Access Control Lists section for more details.

IP-ACL Counter Cleanup

Use the **clear** command to clear the counters for a specified IPv4-ACL filter entry.

Note You cannot use this command to clear the counters for individual filters.

```
switch# show ip access-list abc

ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (2 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (7 matches)

switch# clear ip access-list counters abc
switch# show ip access-list abc

ip access-list abc permit tcp any any (0 matches)
ip access-list abc permit udp any any (0 matches)
ip access-list abc permit icmp any any (0 matches)
ip access-list abc permit ip 10.1.1.0 0.0.0.255 (0 matches)
ip access-list abc permit ip 10.3.70.0 0.0.0.255 (0 matches)
```

Use the **clear ipv6 access-list** command to clear the counters for all IPv6-ACLs.

```
switch# clear ipv6 access-list
```

Use the **clear ipv6 access-list name** command to clear the counters for a specified IPv6-ACL.

```
switch# clear ipv6 access-list List1
```

Note You cannot use this command to clear the counters for each individual filter.