
# Configuring Access Rules

This chapter describes how to control network access through the ASA using access rules and includes the following sections:



**Note** You use access rules to control network access in both routed and transparent firewall modes. In transparent mode, you can use both access rules (for Layer 3 traffic) and EtherType rules (for Layer 2 traffic).

To access the ASA interface for management access, you do not also need an access rule allowing the host IP address. You only need to configure management access according to Chapter 42, "Configuring Management Access."

## Information About Access Rules

Your access policy is made up of one or more access rules and/or EtherType rules per interface or globally for all interfaces.

You can use access rules in routed and transparent firewall mode to control IP traffic. An access rule permits or denies traffic based on the protocol, a source and destination IP address or network, and optionally the source and destination ports.

For transparent mode only, an EtherType rule controls network access for non-IP traffic. An EtherType rule permits or denies traffic based on the EtherType.

This section includes the following topics:

# General Information About Rules

This section describes information for both access rules and EtherType rules, and it includes the following topics:



## Implicit Permits

For routed mode, the following types of traffic are allowed through by default:

- Unicast IPv4 traffic from a higher security interface to a lower security interface.
- Unicast IPv6 traffic from a higher security interface to a lower security interface.

For transparent mode, the following types of traffic are allowed through by default:

- Unicast IPv4 traffic from a higher security interface to a lower security interface.
- Unicast IPv6 traffic from a higher security interface to a lower security interface.
- ARPs in both directions.

**Note** ARP traffic can be controlled by ARP inspection, but cannot be controlled by an access rule.

- BPDUs in both directions.

For other traffic, you need to use either an access rule (IPv4 and IPv6) or an EtherType rule (non-IPv4/IPv6).

## Information About Interface Access Rules and Global Access Rules

You can apply an access rule to a specific interface, or you can apply an access rule globally to all interfaces. You can configure global access rules in conjunction with interface access rules, in which case, the specific interface access rules are always processed before the general global access rules.

**Note** Global access rules apply only to inbound traffic. See the "Inbound and Outbound Rules" section on page 41-3.

## Using Access Rules and EtherType Rules on the Same Interface

You can apply both access rules and EtherType rules to each direction of an interface.

## Rule Order

The order of rules is important. When the ASA decides whether to forward or drop a packet, the ASA tests the packet against each rule in the order in which the rules are listed. After a match is found, no more rules are checked. For example, if you create an access rule at the beginning that explicitly permits all traffic for an interface, no further rules are ever checked. For more information, see the "Implicit Deny" section on page 41-3.

You can disable a rule by making it inactive.

## Implicit Deny

Access lists have an implicit deny at the end of the list, so unless you explicitly permit it, traffic cannot pass. For example, if you want to allow all users to access a network through the ASA except for particular addresses, then you need to deny the particular addresses and then permit all others.

For EtherType access lists, the implicit deny at the end of the access list does not affect IP traffic or ARPs; for example, if you allow EtherType 8037, the implicit deny at the end of the access list does not now block any IP traffic that you previously allowed with an extended access list (or implicitly allowed from a high security interface to a low security interface). However, if you explicitly deny all traffic with an EtherType ACE, then IP and ARP traffic is denied.

If you configure a global access rule, then the implicit deny comes *after* the global rule is processed. See the following order of operations:

1. Interface access rule.
2. Global access rule.
3. Implicit deny.

## Using Remarks

In the ASDM access rule window, a remark that displays next to the rule is the one that was configured before the rule, so when you configure a remark from the CLI and then view it in an ASDM access rule window, the remark displays next to the rule that was configured after the remark in the CLI. However, the packet tracer in ASDM matches the remark that is configured after the matching rule in the CLI.

## NAT and Access Rules

Access rules always use the real IP addresses when determining an access rule match, even if you configure NAT. For example, if you configure NAT for an inside server, 10.1.1.5, so that it has a publicly routable IP address on the outside, 209.165.201.5, then the access rule to allow the outside traffic to access the inside server needs to reference the server's real IP address (10.1.1.5), and not the mapped address (209.165.201.5).

## Inbound and Outbound Rules

The ASA supports two types of access lists:

- Inbound—Inbound access rules apply to traffic as it enters an interface. Global access rules are always inbound.
- Outbound—Outbound access lists apply to traffic as it exits an interface.

**Note** "Inbound" and "outbound" refer to the application of an access list on an interface, either to traffic entering the ASA on an interface or traffic exiting the ASA on an interface. These terms do not refer to the movement of traffic from a lower security interface to a higher security interface, commonly known as inbound, or from a higher to lower interface, commonly known as outbound.

An outbound access list is useful, for example, if you want to allow only certain hosts on the inside networks to access a web server on the outside network. Rather than creating multiple inbound access lists to restrict access, you can create a single outbound access list that allows only the specified hosts. (See Figure 41-1.) The outbound access list prevents any other hosts from reaching the outside network.

***Figure 41-1 Outbound Access List***

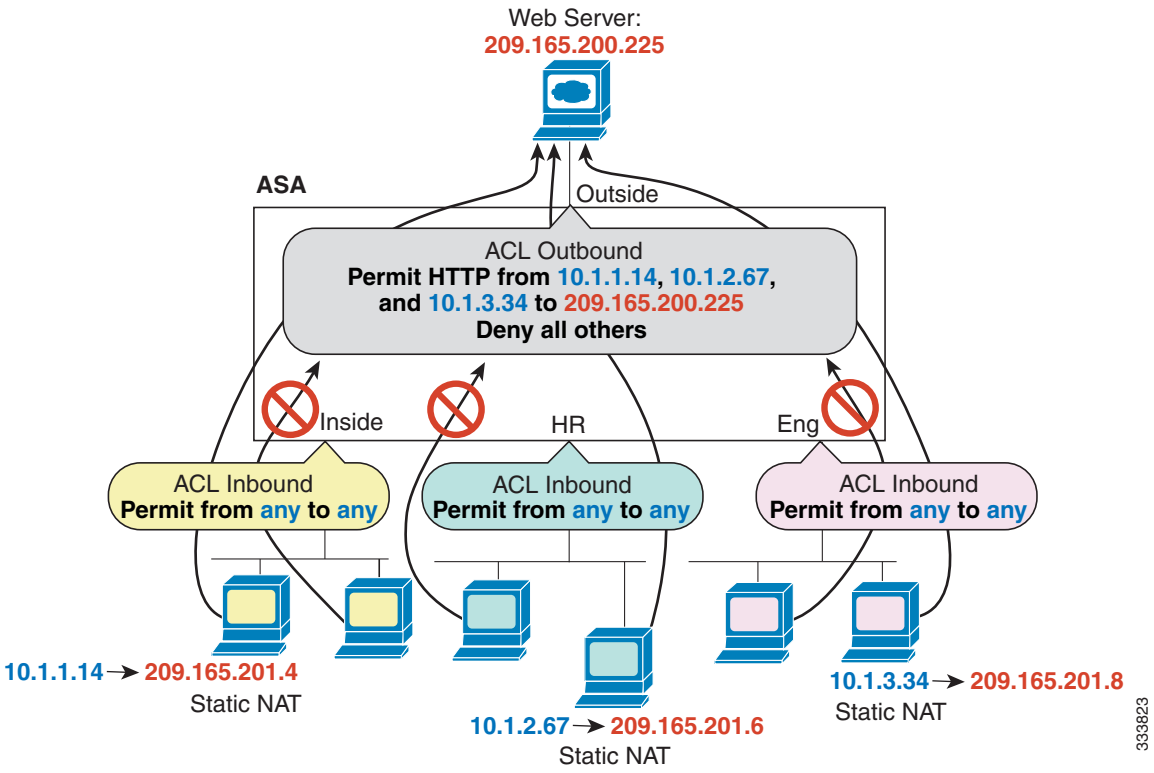


# Information About Access Rules

This section describes information about access rules and includes the following topics:

- Access Rules for Returning Traffic, page 41-5
- Allowing Broadcast and Multicast Traffic through the Transparent Firewall Using Access Rules, page 41-5
- Management Access Rules, page 41-5

## Access Rules for Returning Traffic

For TCP and UDP connections for both routed and transparent mode, you do not need an access rule to allow returning traffic because the ASA allows all returning traffic for established, bidirectional connections.

For connectionless protocols such as ICMP, however, the ASA establishes unidirectional sessions, so you either need access rules to allow ICMP in both directions (by applying access lists to the source and destination interfaces), or you need to enable the ICMP inspection engine. The ICMP inspection engine treats ICMP sessions as bidirectional connections.

## Allowing Broadcast and Multicast Traffic through the Transparent Firewall Using Access Rules

In routed firewall mode, broadcast and multicast traffic is blocked even if you allow it in an access rule, including unsupported dynamic routing protocols and DHCP (unless you configure DHCP relay). Transparent firewall mode can allow any IP traffic through.

**Note** Because these special types of traffic are connectionless, you need to apply an access rule to both interfaces, so returning traffic is allowed through.

Table 41-1 lists common traffic types that you can allow through the transparent firewall.

*Table 41-1 Transparent Firewall Special Traffic*

| Traffic Type | Protocol or Port | Notes |
|---|---|---|
| DHCP | UDP ports 67 and 68 | If you enable the DHCP server, then the ASA does not pass DHCP packets. |
| EIGRP | Protocol 88 | — |
| OSPF | Protocol 89 | — |
| Multicast streams | The UDP ports vary depending on the application. | Multicast streams are always destined to a Class D address (224.0.0.0 to 239.x.x.x). |
| RIP (v1 or v2) | UDP port 520 | — |

## Management Access Rules

You can configure access rules that control management traffic destined to the ASA. Access control rules for to-the-box management traffic (such as HTTP, Telnet, and SSH) have higher precedence than an management access rule. Therefore, such permitted management traffic will be allowed to come in even if explicitly denied by the to-the-box access list.

# Information About EtherType Rules

This section describes EtherType rules and includes the following topics:

### Supported EtherTypes and Other Traffic

An EtherType rule controls the following:

- EtherType identified by a 16-bit hexadecimal number, including common types IPX and MPLS unicast or multicast.
- Ethernet V2 frames.
- BPDUs, which are permitted by default. BPDUs are SNAP-encapsulated, and the ASA is designed to specifically handle BPDUs.
- Trunk port (Cisco proprietary) BPDUs. Trunk BPDUs have VLAN information inside the payload, so the ASA modifies the payload with the outgoing VLAN if you allow BPDUs.
- IS-IS (supported in Version 8.4(5) only).

The following types of traffic are not supported:

- 802.3-formatted frames—These frames are not handled by the rule because they use a length field as opposed to a type field.

### Access Rules for Returning Traffic

Because EtherTypes are connectionless, you need to apply the rule to both interfaces if you want traffic to pass in both directions.

### Allowing MPLS

If you allow MPLS, ensure that Label Distribution Protocol and Tag Distribution Protocol TCP connections are established through the ASA by configuring both MPLS routers connected to the ASA to use the IP address on the ASA interface as the router-id for LDP or TDP sessions. (LDP and TDP allow MPLS routers to negotiate the labels (addresses) used to forward packets.)

On Cisco IOS routers, enter the appropriate command for your protocol, LDP or TDP. The *interface* is the interface connected to the ASA.

```
hostname(config)# mpls ldp router-id interface force
```

Or

```
hostname(config)# tag-switching tdp router-id interface force
```

## Licensing Requirements for Access Rules

| Model | License Requirement |
|---|---|
| All models | Base License. |

## Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

### Context Mode Guidelines

Supported in single and multiple context mode.

### Firewall Mode Guidelines

Supported in routed and transparent firewall modes.

### IPv6 Guidelines

Supports IPv6. (9.0 and later) The source and destination addresses can include any mix of IPv4 and IPv6 addresses. For pre-9.0 versions, you must create a separate IPv6 access rule.

# Default Settings

See the "Implicit Permits" section on page 41-2.

# Configuring Access Rules

This section includes the following topics:



## Adding an Access Rule

To apply an access rule, perform the following steps.

### Detailed Steps

**Step 1** Choose **Configuration > Firewall > Access Rules**.

**Step 2** Click **Add**, and choose one of the following options:

The Add Access Rule dialog box appears.

**Step 3** From the Interface drop-down list, choose the interface on which to apply the rule. Choose **Any** to apply a global rule.

**Step 4** In the Action field, click one of the following radio buttons next to the desired action:

- **Permit**—Permits access if the conditions are matched.
- **Deny**—Denies access if the conditions are matched.

**Step 5** In the Source field, enter an IP address that specifies the network, interface IP, or any address from which traffic is permitted or denied to the specified destination. You may use either an IPv4 or IPv6 address.

For more information about enabling IPv6 on an interface, see the "Configuring IPv6 Addressing" section on page 17-15.

**Step 6** In the User field, enter a user name or group to the access list. Enter the user name in the format *domain_NetBIOS_name\user_name*. Enter the group name in the format *domain_NetBIOS_name\group_name*.

You can configure access rules based on user names and user group names rather than through source IP addresses. The ASA applies the security policies based on an association of IP addresses to Windows Active Directory login information and reports events based on the mapped user names instead of network IP addresses.

See the "Configuring Identity-Based Security Policy" section on page 47-15 for more information.

**Step 7** To browse for a user name or user group, click the ellipsis (...) button. The Browse User dialog box appears.

**Step 8** In the Destination field, enter an IP address that specifies the network, interface IP, any address to which traffic is permitted or denied from the source specified in the Source field. You may use either an IPv4 or IPv6 address.

**Step 9** Select the service type.

**Step 10** (Optional) To add a time range to your access rule that specifies when traffic can be allowed or denied, click **More Options** to expand the list.

a. To the right of the Time Range drop down list, click the browse button.

   The Browse Time Range dialog box appears.

b. Click **Add**.

   The Add Time Range dialog box appears.

c. In the Time Range Name field, enter a time range name, with no spaces.

d. Choose the Start Time and the End Time.

e. To specify additional time constraints for the time range, such as specifying the days of the week or the recurring weekly interval in which the time range will be active, click **Add**, and choose the specifications.

f. Click **OK** to apply the optional time range specifications.

**Step 11** (Optional) In the Description field, add a text description about the access rule.

The description can contain multiple lines; however, each line can be no more than 100 characters in length.

**Step 12** (Optional) Logging is enabled by default. You can disable logging by unchecking the check box, or you can change the logging level from the drop-down list. The default logging level is Informational.

**Step 13** Click **OK**. The access rule appears with the newly configured access rules.

**Step 14** Click **Apply** to save the access rule to your configuration.

You can edit or delete a particular access rule by selecting the rule and then clicking Edit or Delete.

# Adding an EtherType Rule (Transparent Mode Only)

The EtherType Rules window shows access rules based on packet EtherTypes. EtherType rules are used to configure non-IP related traffic policies through the ASA when operating in transparent mode. In transparent mode, you can apply both extended and EtherType access rules to an interface. EtherType rules take precedence over the extended access rules.

For more information about EtherType rules, see the "Information About Access Rules" section on page 41-1.

To add an EtherType rule, perform the following steps:

**Step 1** Choose **Configuration > Device Management > Management Access > EtherType Rules**.

**Step 2** Click **Add**.

The Add EtherType rules window appears.

**Step 3** (Optional) To specify the placement of the new EtherType rule, select an existing rule, and click **Insert...** to add the EtherType rule before the selected rule, or click **Insert After...** to add the EtherType rule after the selected rule.

**Step 4** From the Interface drop-down list, choose the interface on which to apply the rule. Choose **Any** to apply a global rule.

**Step 5** In the Action field, click one of the following radio buttons next to the desired action:

- **Permit**—Permits access if the conditions are matched.
- **Deny**—Denies access if the conditions are matched.

**Step 6** In the EtherType field, choose an EtherType value from the drop-down list.

**Step 7** (Optional) In the Description field, add a test description about the rule.

The description can contain multiple lines; however, each line can be no more than 100 characters in length.

**Step 8** (Optional) To specify the direction for this rule, click **More Options** to expand the list, and then specify the direction by clicking one of the following radio buttons:

- **In**—Incoming traffic
- **Out**—Outgoing traffic

**Step 9** Click **OK**.

# Configuring Management Access Rules

You can configure an interface ACL that supports access control for to-the-box management traffic from a specific peer (or set of peers) to the security appliance. One scenario in which this type of ACL would be useful is when you want to block IKE Denial of Service attacks.

To configure an extended ACL that permits or denies packets for to-the-box traffic, perform the following steps:

**Step 1** Choose **Configuration > Device Management > Management Access > Management Access Rules**.

**Step 2** Click **Add**, and choose one of the following actions:

The Add Management Access Rule dialog box appears.

**Step 3** From the Interface drop-down list, choose an interface on which to apply the rule. Choose **Any** to apply a global rule.

**Step 4** In the Action field, click one of the following radio buttons to choose the action:

- **Permit**—Permits access if the conditions are matched.

- **Deny**—Denies access if the conditions are matched.

**Step 5** In the Source field, enter an IP address that specifies the network object group, interface IP, or any address from which traffic is permitted or denied. You may use either an IPv4 or IPv6 address.

**Note** IPv6 must be enabled on at least one interface before you can configure an extended ACL with an IPv6 address. For more information about enabling IPv6 on an interface, see the "Configuring IPv6 Addressing" section on page 17-15

**Step 6** In the Service field, add a service name for rule traffic, or click the ellipsis (...) to browse for a service.

**Step 7** (Optional) In the Description field, add a description for this management access rule.

The description can contain multiple lines; however, each line can be no more than 100 characters in length.

**Step 8** (Optional) Logging is enabled by default. You can disable logging by unchecking the check box, or you can change the logging level from the drop-down list. The default logging level is Informational.

**Step 9** (Optional) To add a source service (TCP, UDP, and TCP-UDP only) and a time range to your access rule that specifies when traffic can be allowed or denied, click **More Options** to expand the list.If you want to turn off this Management Access Rule, uncheck **Enable Rule**.

- Add a source service in the Source Service field, or click the ellipsis (...) to browse for a service.

  The destination service and source service must be the same. Copy and paste the destination Service field to the Source Service field.
- To configure the logging interval (if you enable logging and choose a non-default setting), enter a value in seconds in the Logging Interval field.
- To select a predefined time range for this rule, from the Time Range drop-down list, choose a time range; or click the ellipsis (...) to browse for a time range. You can also specify additional time constraints for the time range, such as specifying the days of the week or the recurring weekly interval in which the time range will be active.

**Step 10** Click **OK**. The dialog box closes, and the Management Access rule is added.

**Step 11** Click **Apply**. The rule is saved in the running configuration.

# Advanced Access Rule Configuration

The Advanced Access Rule Configuration dialog box lets you to set global access rule logging options.

When you enable logging, if a packet matches the access rule, the ASA creates a flow entry to track the number of packets received within a specific interval. The ASA generates a system log message at the first hit and at the end of each interval, identifying the total number of hits during the interval and reporting the time of the last hit.

**Note** The ASApane displays the hit count information in the "last rule hit" row. To view the rule hit count and timestamp, choose **Configuration > Firewall > Advanced > ACL Manager**, and hover the mouse pointer over a cell in the ACL Manager table.

At the end of each interval, the ASA resets the hit count to 0. If no packets match the access rule during an interval, the ASA deletes the flow entry.

A large number of flows can exist concurrently at any point of time. To prevent unlimited consumption of memory and CPU resources, the ASA places a limit on the number of concurrent deny flows; the limit is placed only on deny flows (and not permit flows) because they can indicate an attack. When the limit is reached, the ASA does not create a new deny flow until the existing flows expire. If someone initiates a denial of service attack, the ASA can create a very large number of deny flows in a very short period of time. Restricting the number of deny-flows prevents unlimited consumption of memory and CPU resources.

#### Prerequisites

These settings only apply if you enable the newer logging mechanism for the access rule.

#### Fields

- Maximum Deny-flows—The maximum number of deny flows permitted before the ASA stops logging, between 1 and the default value. The default is 4096.
- Alert Interval—The amount of time (1-3600 seconds) between system log messages (number 106101) that identify that the maximum number of deny flows was reached. The default is 300 seconds.
- Per User Override table—Specifies the state of the per user override feature. If the per user override feature is enabled on the inbound access rule, the access rule provided by a RADIUS server replaces the access rule configured on that interface. If the per user override feature is disabled, the access rule provided by the RADIUS server is combined with the access rule configured on that interface. If the inbound access rule is not configured for the interface, per user override cannot be configured.
- 

  For VPN remote access traffic, the behavior depends on whether there is a VPN filter applied in the group policy (see the Configuration > Remote Access VPN > Network (Client) Access > Group Policies > Add/Edit > General > More Options > Filter field) and whether you set the Per User Override option:

  – No Per User Override, no VPN filter —Traffic is matched against the interface ACL (per the default **Enable inbound VPN sessions to bypass interface access lists** setting (disabled) on the Configurartion > Remote Access VPN > Network (Client) Access > AnyConnect Connection Profiles pane).

  – No Per User Override, VPN filter —Traffic is matched first against the interface ACL, then against the VPN filter.

  – Per User Override, VPN filter —Traffic is matched against the VPN filter only.
- Object Group Search Setting—Reduces the amount of memory used to store service rules, but lengthens the amount of time to search for a matching access rule.

## Access Rule Explosion

The security appliance allows you to turn off the expansion of access rules that contain certain object groups. When expansion is turned off, an object group search is used for lookup, which lowers the memory requirements for storing expanded rules but decreases the lookup performance. Because of the trade-off of performance for memory utilization, you can turn on and turn off the search.

To configure the option of turning off the expansion of access rules that contain s, perform the following steps:

**Step 1** Choose **Configuration > Firewall > Access Rules**.

**Step 2** Click the **Advanced** button.

**Step 3** Check the **Enable Object Group Search Algorithm** check box.

# Configuring HTTP Redirect

The HTTP Redirect table displays each interface on the ASA, shows whether it is configured to redirect HTTP connections to HTTPS, and the port number from which it redirects those connections.

**Note** To redirect HTTP, the interface requires an access list that permits HTTP. Otherwise, the interface cannot listen to the HTTP port.

The Configuration > Device Management > Advanced > HTTP Redirect > Edit pane lets you change the HTTP redirect setting of an interface or the port from which it redirects HTTP connections. Select the interface in the table and click **Edit**. You can also double-click an interface. The Edit HTTP/HTTPS Settings dialog box opens.

## Edit HTTP/HTTPS Settings

The Edit HTTP/HTTPS Settings dialog box lets you change the HTTP redirect setting of an interface or the port number.

### Fields

The Edit HTTP/HTTPS Settings dialog box includes the following fields:

- Interface—Identifies the interface on which the ASA redirects or does not redirect HTTP requests to HTTPS.
- Redirect HTTP to HTTPS—Check to redirect HTTP requests to HTTPS, or uncheck to not redirect HTTP requests to HTTPS.
- HTTP Port—Identifies the port from which the interface redirects HTTP connections. By default it listens to port 80.

For more information about access rules, see the "Information About Access Rules" section on page 41-1.

# Feature History for Access Rules

Table 41-2 lists each feature change and the platform release in which it was implemented. ASDM is backwards-compatible with multiple platform releases, so the specific ASDM release in which support was added is not listed.

***Table 41-2 Feature History for Access Rules***

| Feature Name | Platform Releases | Feature Information |
|---|---|---|
| Interface access rules | 7.0(1) | Controlling network access through the ASA using access lists.<br><br>We introduced the following screen: Configuration > Firewall > Access Rules. |
| Global access rules | 8.3(1) | Global access rules were introduced.<br><br>We modified the following screen: Configuration > Firewall > Access Rules. |
| Support for Identity Firewall | 8.4(2) | You can now use identity firewall users and groups for the source and destination. You can use an identity firewall ACL with access rules, AAA rules, and for VPN authentication. |
| EtherType ACL support for IS-IS traffic (transparent firewall mode) | 8.4(5) | In transparent firewall mode, the ASA can now pass IS-IS traffic using an EtherType ACL.<br><br>We modified the following screen: Configuration > Device Management > Management Access > EtherType Rules.<br><br>*Not available in Version 8.5(1), 8.6(1), or 9.0(1).* |
| Support for TrustSec | 9.0(1) | You can now use TrustSec security groups for the source and destination. You can use an identity firewall ACL with access rules. |

*Table 41-2 Feature History for Access Rules (continued)*

| Feature Name | Platform Releases | Feature Information |
|---|---|---|
| Unified ACL for IPv4 and IPv6 | 9.0(1) | ACLs now support IPv4 and IPv6 addresses. You can even specify a mix of IPv4 and IPv6 addresses for the source and destination. The **any** keyword was changed to represent IPv4 and IPv6 traffic. The **any4** and **any6** keywords were added to represent IPv4-only and IPv6-only traffic, respectively. The IPv6-specific ACLs are deprecated. Existing IPv6 ACLs are migrated to extended ACLs. See the release notes for more information about migration.<br>We modified the following screens:<br>Configuration > Firewall > Access Rules<br>Configuration > Remote Access VPN > Network (Client) Access > Group Policies > General > More Options |
| Extended ACLand object enhancement to filter ICMP traffic by ICMP code | 9.0(1) | ICMP traffic can now be permitted/denied based on ICMP code.<br>We introduced or modified the following screens:<br>Configuration > Firewall > Objects > Service Objects/Groups<br>Configuration > Firewall > Access Rule |