

Access Rules

This chapter describes how to control network access through or to the ASA using access rules. 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).

Note 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 the general operations configuration guide.

- Controlling Network Access, on page 1
- Licensing for Access Rules, on page 7
- Guidelines for Access Control, on page 7
- Configure Access Control, on page 8
- Monitoring Access Rules, on page 17
- History for Access Rules, on page 18

Controlling Network Access

Access rules determine which traffic is allowed through the ASA. There are several different layers of rules that work together to implement your access control policy:

- Extended access rules (Layer 3+ traffic) assigned to interfaces—You can apply separate rule sets (ACLs) in the inbound and outbound directions. An extended access rule permits or denies traffic based on the source and destination traffic criteria.
- Extended access rules (Layer 3+ traffic) assigned to Bridge Virtual Interfaces (BVI; routed mode)—If you name a BVI, you can apply separate rule sets in the inbound and outbound direction, and you can also apply rule sets to the bridge group member interfaces. When both the BVI and member interface have access rules, the order of processing depends on direction. Inbound, the member access rules are evaluated first, then the BVI access rules. Outbound, the BVI rules are considered first, then the member interface rules.
- Extended access rules assigned globally—You can create a single global rule set, which serves as your default access control. The global rules are applied after interface rules.

- Management access rules (Layer 3+ traffic)—You can apply a single rule set to cover traffic directed at an interface, which would typically be management traffic. In the CLI, these are “control plane” access groups. For ICMP traffic directed at the device, you can alternatively configure ICMP rules.
- EtherType rules (Layer 2 traffic) assigned to interfaces (bridge group member interfaces only)—You can apply separate rule sets in the inbound and outbound directions. EtherType rules control network access for non-IP traffic. An EtherType rule permits or denies traffic based on the EtherType. You can also apply extended access rules to bridge group member interfaces to control Layer 3+ traffic.

General Information About Rules

The following topics provide general information about access rules and EtherType rules.

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 inbound interface access rules are always processed before the general global access rules. Global access rules apply only to inbound traffic.

Inbound and Outbound Rules

You can configure access rules based on the direction of traffic:

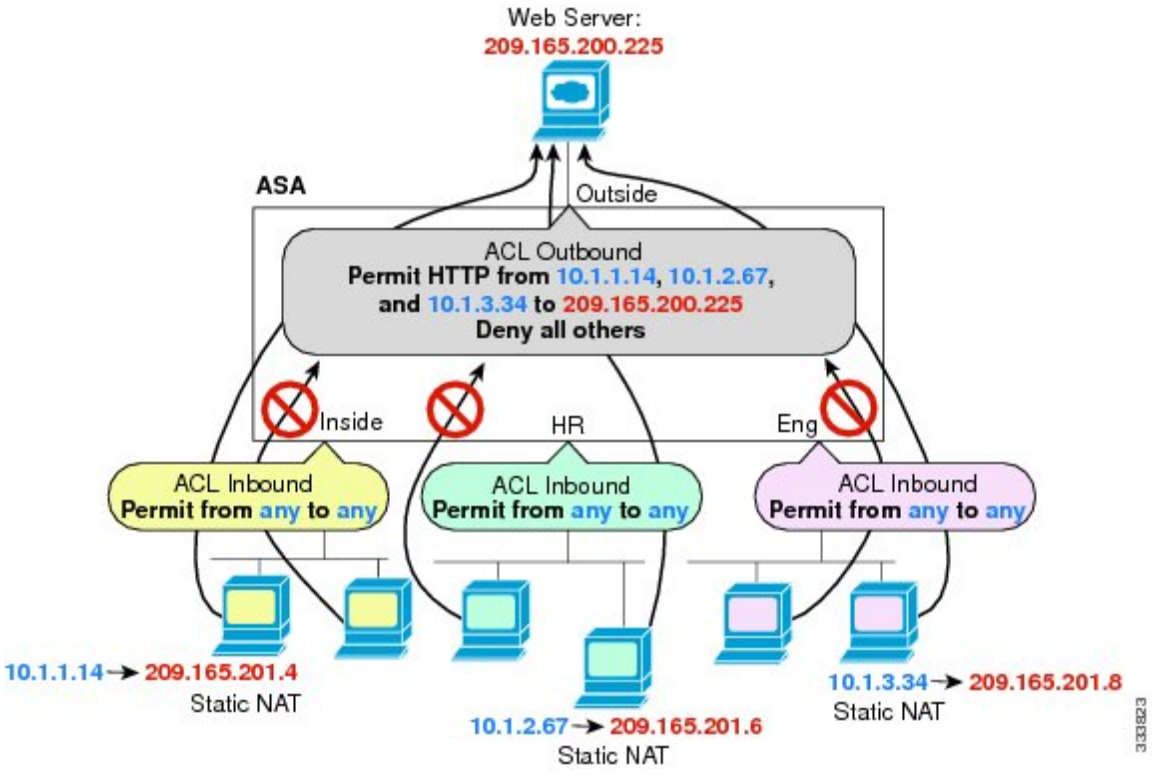
- Inbound—Inbound access rules apply to traffic as it enters an interface. Global and management access rules are always inbound.
- Outbound—Outbound rules apply to traffic as it exits an interface.



Note “Inbound” and “outbound” refer to the application of an ACL 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 ACL 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 ACLs to restrict access, you can create a single outbound ACL that allows only the specified hosts. (See the following figure.) The outbound ACL prevents any other hosts from reaching the outside network.

Figure 1: Outbound ACL



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 in the applied ACL. 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.

Implicit Permits

Unicast IPv4 and IPv6 traffic from a higher security interface to a lower security interface is allowed through by default. This includes traffic between standard routed interfaces and Bridge Virtual Interfaces (BVI) in routed mode.

For bridge group member interfaces, this implicit permit from a higher to a lower security interface applies to interfaces within the same bridge group only. There are no implicit permits between a bridge group member interface and a routed interface or a member of a different bridge group.

Bridge group member interfaces (routed or transparent mode) also allow the following by default:

- ARPs in both directions. (You can control ARP traffic using ARP inspection, but you cannot control it by access rule.)
- BPDUs in both directions. (You can control these using Ethertype rules.)

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

Implicit Deny

ACLs 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 management (control plane) ACLs, which control to-the-box traffic, there is no implicit deny at the end of a set of management rules for an interface. Instead, any connection that does not match a management access rule is then evaluated by regular access control rules.

For EtherType ACLs, the implicit deny at the end of the ACL does not affect IP traffic or ARPs; for example, if you allow EtherType 8037, the implicit deny at the end of the ACL does not now block any IP traffic that you previously allowed with an extended ACL (or implicitly allowed from a high security interface to a low security interface). However, if you explicitly deny all traffic with an EtherType rule, then IP and ARP traffic is denied; only physical protocol traffic, such as auto-negotiation, is still allowed.

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. For bridge group member interfaces, the Bridge Virtual Interface (BVI) access rule.
3. Global access rule.
4. Implicit deny.

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).

Same Security Level Interfaces and Access Rules

Each interface has a security level, and security level checking is performed before access rules are considered. Thus, even if you allow a connection in an access rule, it can be blocked due to same-security-level checking at the interface level. You might want to ensure that your configuration allows same-security-level connections so that your access rules are always considered for permit/deny decisions.

- Connections between the same security level ingress and egress interfaces are subject to the same-security-traffic inter-interface check.

  To allow these connections, enter the **same-security-traffic permit inter-interface** command.

  To allow these connections, choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, then select the **Enable traffic between two or more interfaces which are configured with the same security levels** option.

- Connections with the same ingress and egress interfaces are subject to the same-security-traffic intra-interface check.

  To allow these connections, enter the **same-security-traffic permit intra-interface** command.

  To allow these connections, choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, then select the **Enable traffic between two or more hosts connected to the same interface** option.

Extended Access Rules

This section describes information about extended access rules.

Extended Access Rules for Returning Traffic

For TCP, UDP, and SCTP 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 ACLs 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. For example, to control ping, specify **echo-reply** (**0**) (ASA to host) or **echo** (**8**) (host to ASA).

Allowing Broadcast and Multicast Traffic

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. You must configure the dynamic routing protocols or DHCP relay to allow this traffic.

For interfaces that are members of the same bridge group in transparent or routed firewall mode, you can allow any IP traffic through using access rules.

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

The following table lists common traffic types that you can allow using access rules between interfaces that are members of the same bridge group.

Table 1: Special Traffic for Access Rules between Members of the Same Bridge Group

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 connections to an interface) have higher

precedence than a management access rule . Therefore, such permitted management traffic will be allowed to come in even if explicitly denied by the to-the-box ACL.

Unlike regular access rules, there is no implicit deny at the end of a set of management rules for an interface. Instead, any connection that does not match a management access rule is then evaluated by regular access control rules.

Alternatively, you can use ICMP rules to control ICMP traffic to the device. Use regular extended access rules to control ICMP traffic through the device.

EtherType Rules

This section describes EtherType rules.

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.
- Intermediate System to Intermediate System (IS-IS).
- The IEEE 802.2 Logical Link Control packet. You can control access based on the Destination Service Access Point address.

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.

EtherType 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.

mpls ldp router-id *interface* **force**

Or

tag-switching tdp router-id *interface* **force**

Licensing for Access Rules

Access control rules do not require a special license.

However, to use **sctp** as the protocol in a rule, you must have a Carrier license.

Guidelines for Access Control

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.

Per-User ACL Guidelines

- The per-user ACL uses the value in the **timeout uauth** command, but it can be overridden by the AAA per-user session timeout value.
- If traffic is denied because of a per-user ACL, syslog message 109025 is logged. If traffic is permitted, no syslog message is generated. The **log** option in the per-user ACL has no effect.

Additional Guidelines and Limitations

- Over time, your list of access rules can grow to include many obsolete rules. Eventually, the ACLs for the access groups can become so large that they impact overall system performance. If you find that the system is having issues sending syslog messages, communicating for failover synchronization, establishing and maintaining SSH/HTTPS management access connections, and so forth, you might need to prune your access rules. In general, you should actively maintain your rule lists to remove obsolete rules, rules that are never hit, FQDN objects that can no longer be resolved, and so forth. Also consider implementing object group search.
- Object group search is enabled by default for new deployments.

  You can reduce the memory required to search access rules by enabling object group search, but this is at the expense rule of lookup performance and increased CPU utilization. When enabled, object group search does not expand network or service objects, but instead searches access rules for matches based on those group definitions. You can set this option by clicking the **Advanced** button below the access rule table.

  You can use the **object-group-search threshold** command to enable a threshold to help prevent performance degradation. When operating with a threshold, for each connection, both the source and destination IP addresses are matched against network objects. If the number of objects matched by the source address times the number matched by the destination address exceeds 10,000, the connection is dropped. Configure your rules to prevent an excessive number of matches.

Note Object group search works with network and service objects only. It does not work with security group or user objects. Do not enable this feature if the ACLs include security groups. The result can be inactive ACLs or other unexpected behavior.

- You can improve system performance and reliability by using the transactional commit model for access groups. However, transactional commit is not recommended if you use FQDN objects for hostnames whose resolution can frequently change, as the access group compilation might never completely resolve. See the basic settings chapter in the general operations configuration guide for more information. The option is under **Configurations > Device Management > Advanced > Rule Engine**.
- In ASDM, rule descriptions are based on the access list remarks that come before the rule in the ACL; for new rules you create in ASDM, any descriptions are also configured as remarks before the related rule. However, the packet tracer in ASDM matches the remark that is configured after the matching rule in the CLI.
- If you enter more than one item in source or destination address, or source or destination service, ASDM automatically creates an object group for them with the prefix DM_INLINE. These objects are automatically expanded to their component parts in the rule table view, but you can see the object names if you deselect the **Auto-expand network and service objects with specified prefix** rule table preference in **Tools** > **Preferences**.
- To use fully-qualified domain name (FQDN) network objects as source or destination criteria, you must also configure DNS for the data interfaces.

  Note that controlling access by FQDN is a best-effort mechanism. Consider the following points:

  - Because DNS replies can be spoofed, only use fully trusted internal DNS servers.
  - Some FQDNs, especially for very popular servers, can have hundreds if not thousands of IP addresses, and these can frequently change. Because the system uses cached DNS lookup results, users might get addresses that are not yet in the cache, and their connections will not match the FQDN rule. Rules that use FQDN network objects function effectively only for names that resolve to fewer than 100 addresses.

    We recommend that you do not create network object rules for an FQDN that resolves to more than 100 addresses, as the likelihood of the address in a connection being one that has been resolved and available in the DNS cache on the device is low.
  - For popular FQDNs, different DNS servers can return a different set of IP addresses. Thus, if your users use a different DNS server than the one you configure, FQDN-based access control rules might not apply to all IP addresses for the site that are used by your clients, and you will not get the intended results for your rules.
  - Some FQDN DNS entries have very short time to live (TTL) values. This can result in frequent recompliation of the lookup table, which can impact overall system performance.

Configure Access Control

The following topics explain how to configure access control.

Configure Access Rules

To apply an access rule, perform the following steps.

Procedure

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

The rules are organized by interface and direction, with a separate group for global rules. If you configure management access rules, they are repeated on this page. These groups are equivalent to the extended ACL that is created and assigned to the interface or globally as an access group. These ACLs also appear on the ACL Manager page.

Step 2 Do any of the following:

- To add a new rule, choose **Add > Add Access Rule**.
- To insert a rule at a specific location within a container, select an existing rule and choose **Add > Insert** to add the rule above it, or choose **Add > Insert After**.
- To edit a rule, select it and click **Edit**.

Step 3 Fill in the rule properties. The primary options to select are:

- **Interface**—The interface to which the rule applies. Select Any to create a global rule. For bridge groups in routed mode, you can create access rules for both the Bridge Virtual Interface (BVI) and each bridge group member interface.
- **Action: Permit/Deny**—Whether you are permitting (allowing) the described traffic or are denying (dropping) it.
- **Source/Destination criteria**—A definition of the source (originating address) and destination (target address of the traffic flow). You typically configure IPv4 or IPv6 addresses of hosts or subnets, which you can represent with network or network object groups, or network-service object groups. You can also specify a user or user group name for the source. Additionally, you can use the Service field to identify the specific type of traffic if you want to focus the rule more narrowly than all IP traffic. If you include service specifications in a network-service object, specify IP in the Service field. If you implement Trustsec, you can use security groups to define source and destination.

For detailed information on all of the available options, see Access Rule Properties, on page 9.

When you are finished defining the rule, click **OK** to add the rule to the table.

Step 4 Click **Apply** to save the access rule to your configuration.

Access Rule Properties

When you add or edit an access rule, you can configure the following properties. In many fields, you can click the "..." button on the right of the edit box to select, create, or edit objects that are available for the field.

Interface

The interface to which the rule applies. Select Any to create a global rule. For bridge groups in routed mode, you can select either the Bridge Virtual Interface (BVI) or the bridge group member interfaces.

Action: Permit/Deny

Whether you are permitting (allowing) the described traffic or are denying (dropping) it.

Source Criteria

The characteristics of the originator of the traffic you are trying to match. You must configure Source, but the other properties are optional.

Source

The IPv4 or IPv6 address of the source. The default is **any**, which matches all IPv4 or IPv6 addresses; you can use **any4** to target IPv4 only, or **any6** to target IPv6 only. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), a subnet (in 10.100.10.0/24 or 10.100.10.0/255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60), the name of a network object or network object group, the name of a network-service object group, or the name of an interface.

User

If you enable the identity firewall, you can specify a user or user group as the traffic source. The IP address the user is currently using will match the rule. You can specify a username (DOMAIN\user), a user group (DOMAIN\\group, note the double \ indicates a group name), or a user object group. For this field, it is far easier to click "**...**" to select names from your AAA server group than to type them in.

Security Group

If you enable Cisco Trustsec, you can specify a security group name or tag (1-65533), or security group object.

More Options > Source Service

If you specify TCP, UDP, or SCTP as the destination service, you can optionally specify a predefined service object for TCP, UDP, TCP-UDP, or SCTP, or use your own object. Typically, you define the destination service only and not the source service. Note that if you define the source service, the destination service protocol must match it (for example, both TCP, with or without port definitions).

Destination Criteria

The characteristics of the target of the traffic you are trying to match. You must configure Destination, but the other properties are optional.

Destination

The IPv4 or IPv6 address of the destination. The default is **any**, which matches all IPv4 or IPv6 addresses; you can use **any4** to target IPv4 only, or **any6** to target IPv6 only. You can specify a single host address (such as 10.100.10.5 or 2001:DB8::0DB8:800:200C:417A), a subnet (in 10.100.10.0/24 or 10.100.10.0/255.255.255.0 format, or for IPv6, 2001:DB8:0:CD30::/60), the name of a network object or network object group, the name of a network-service object group, or the name of an interface.

Security Group

If you enable Cisco Trustsec, you can specify a security group name or tag (1-65533), or security group object.

Service

The protocol of the traffic, such as IP, TCP, UDP, and optionally ports for TCP, UDP, or SCTP. The default is IP, but you can select a more specific protocol to target traffic with more granularity. Typically, you would select some type of service object. For TCP, UDP, and SCTP, you can specify ports, for example, tcp/80, tcp/http, tcp/10-20 (for a range of ports), tcp-udp/80 (match any TCP or UDP traffic on port 80), sctp/diameter, and so forth. If you include service specifications in a network-service object, specify IP in the Service field.

Description

A explanation of the purpose of the rule, up to 100 characters per line. You can enter multiple lines; each line is added as a remark in the CLI, and the remarks are placed before the rule.



Note If you add remarks with non-English characters on one platform (such as Windows) then try to remove them from another platform (such as Linux), you might not be able to edit or delete them because the original characters might not be correctly recognized. This limitation is due to an underlying platform dependency that encodes different language characters in different ways.

Enable Logging; Logging Level; More Options > Logging Interval

The logging options define how syslog messages will be generated for rules. You can implement the following logging options:

Deselect Enable Logging

This will disable logging for the rule. No syslog messages of any type will be issued for connections that match this rule.

Select Enable Logging with Logging Level = Default

This provides the default logging for rules. Syslog message 106023 is issued for each denied connection. If the appliance comes under attack, the frequency of issuing this message could impact services.

Select Enable Logging with Non-Default Logging Level

This provides a summarized syslog message, 106100, instead of 106023. Message 106100 is issued upon first hit, then again at each interval configured in **More Options > Logging Interval** (default is every 300 seconds, you can specify 1-600), showing the number of hits during that interval. The recommended logging level is **Informational**.

Summarizing deny messages can reduce the impact of attacks and possibly make it easier for you to analyze messages. If you do come under a denial of service attack, you might see message 106101, which indicates that the number of cached deny flows used to produce the hit count for message 106100 has exceeded the maximum for an interval. At this point, the appliance stops collecting statistics until the next interval to mitigate the attack.

More Options > Traffic Direction

Whether the rule is for the **In** or **Out** direction. **In** is the default, and it is the only option for global and management access rules.

More Options > Enable Rule

Whether the rule is active on the device. Disabled rules appear with strike-through text in the rule table. Disabling a rule lets you stop its application to traffic without deleting it, so you can enable it again later if you decide you need it.

More Options > Time Range

The name of the time range object that defines the times of day and days of the week when the rule should be active. If you do not specify a time range, the rule is always active.

Configure Advanced Options for Access Rules

Advanced access rule options allow you to customize certain aspects of rule behavior, but these options have defaults that are appropriate in most cases.

Procedure

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

Step 2 Click the **Advanced** button below the rule table.

Step 3 Configure the following options as required:

- **Advanced Logging Settings**—If you configure non-default logging, the system caches deny flows to develop statistics for message 106100, as explained in Evaluating Syslog Messages for Access Rules, on page 17. To prevent unlimited consumption of memory and CPU resources, the ASA places a limit on the number of concurrent *deny* flows because they can indicate an attack. Message 106101 is issued when the limit is reached. You can control the following aspects related to 106101.
  - Maximum Deny-flows—The maximum number of deny flows permitted before the ASA stops caching flows, between 1 and 4096. The default is 4096.
  - Alert Interval—The amount of time (1-3600 seconds) between issuing system log message 106101, which indicates that the maximum number of deny flows was reached. The default is 300 seconds.

- **Per User Override table**—Whether to allow a dynamic user ACL that is downloaded for user authorization from a RADIUS server to override the ACL assigned to the interface. For example, if the interface ACL denies all traffic from 10.0.0.0, but the dynamic ACL permits all traffic from 10.0.0.0, then the dynamic ACL overrides the interface ACL for that user. Check the **Per User Override** box for each interface that should allow user overrides (inbound direction only). 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.

  By default, VPN remote access traffic is not matched against interface ACLs. However, if you deselect the **Enable inbound VPN sessions to bypass interface access lists** setting on the Configuration > Remote Access VPN > Network (Client) Access > AnyConnect Client Connection Profiles pane), 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.
  - 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**—You can reduce the memory required to search access rules that use object groups by selecting **Enable Object Group Search Algorithm**, but this is at the expense of rule lookup performance. When enabled, object group search does not expand network objects, but instead searches access rules for matches based on those group definitions.

  Select **Enable Object Group Search Threshold** to set a threshold to help prevent performance degradation. When operating with a threshold, for each connection, both the source and destination IP addresses are matched against network objects. If the number of objects matched by the source address times the number matched by the destination address exceeds 10,000, the connection is dropped. Configure your rules to prevent an excessive number of matches.

  Note
  Object group search works with network and service objects only. It does not work with security group objects. Do not enable this feature if the ACLs include security groups. The result can be inactive ACLs or other unexpected behavior.

- **Forward Reference Setting**—(Pre 7.18 only.)Normally, you cannot reference an object or object group that does not exist in an ACL or object group, or delete one that is currently referenced. You also cannot reference an ACL that does not exist in an **access-group** command (to apply access rules). However, you can change this default behavior so that you can "forward reference" objects or ACLs before you create them. Until you create the objects or ACLs, any rules or access groups that reference them are ignored. Select **Enable the forward reference of objects and object-groups** to enable forward referencing. Be aware that if you enable forward referencing, ASDM cannot tell the difference between a typo reference to an existing object and a forward reference.

  Note
  This setting is enabled by default and is no longer configurable starting with ASA 9.18(1).

Step 4 Click **OK**.

Configure Management Access Rules

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

Unlike regular access rules, there is no implicit deny at the end of a set of management rules for an interface. Instead, any connection that does not match a management access rule is then evaluated by regular access control rules.

Procedure

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

The rules are organized by interface. Each group is equivalent to the extended ACL that is created and assigned to the interface as a control plane ACL. These ACLs also appear on the Access Rules and ACL Manager pages.

Step 2 Do any of the following:

- To add a new rule, choose **Add > Add Management Access Rule**.
- To insert a rule at a specific location within a container, select an existing rule and choose **Add > Insert** to add the rule above it, or choose **Add > Insert After**.
- To edit a rule, select it and click **Edit**.

Step 3 Fill in the rule properties. The primary options to select are:

- **Interface**—The interface to which the rule applies. For bridge groups in routed mode, you can create access rules for both the Bridge Virtual Interface (BVI) and each bridge group member interface.
- **Action: Permit/Deny**—Whether you are permitting (allowing) the described traffic or are denying (dropping) it.
- **Source/Destination criteria**—A definition of the source (originating address) and destination (target address of the traffic flow). You typically configure IPv4 or IPv6 addresses of hosts or subnets, which you can represent with network or network object groups. You can also specify a user or user group name for the source. Additionally, you can use the Service field to identify the specific type of traffic if you want to focus the rule more narrowly than all IP traffic. If you implement Trustsec, you can use security groups to define source and destination.

For detailed information on all of the available options, see Access Rule Properties, on page 9.

When you are finished defining the rule, click **OK** to add the rule to the table.

Step 4 Click **Apply** to save the rule to your configuration.

Configure EtherType Rules

EtherType rules apply to non-IP layer-2 traffic on bridge group member interfaces (in routed or transparent firewall mode). You can use these rules to permit or drop traffic based on the EtherType value in the layer-2 packet. With EtherType rules, you can control the flow of non-IP traffic across the ASA.

You can apply both extended and EtherType access rules to a bridge group member interface. EtherType rules take precedence over the extended access rules.

Procedure

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

The rules are organized by interface and direction. Each group is equivalent to the EtherType ACL that is created and assigned to the interface.

Step 2 Do any of the following:

- To add a new rule, choose **Add** > **Add EtherType Rule**.
- To insert a rule at a specific location within a container, select an existing rule and choose **Add** > **Insert** to add the rule above it, or choose **Add** > **Insert After**.

- To edit a rule, select it and click **Edit**.

Step 3 Fill in the rule properties. The primary options to select are:

- **Interface**—The interface to which the rule applies.
- **Action: Permit/Deny**—Whether you are permitting (allowing) the described traffic or are denying (dropping) it.
- EtherType—You can match traffic using the following options:
  - **any**—Matches all traffic.
  - **bpdu**—Bridge protocol data units, which are allowed by default. When you apply the configuration, this keyword is converted to **dsap bpdu** on the device.
  - **dsap**—The IEEE 802.2 Logical Link Control packet's Destination Service Access Point address. You also need to include the address you want to permit or deny in hexadecimal, from 0x01 to 0xff, in **DSAP Value**. Following are the values for some common addresses:
    - **0x42**—bridge protocol data units (BPDU). When you apply the configuration, this is converted to **dsap bpdu** on the device.
    - **0xe0**—Internet Packet Exchange (IPX) 802.2 LLC. When you apply the configuration, this is converted to **dsap ipx** on the device.
    - **0xfe**—Intermediate System to Intermediate System (IS-IS). When you apply the configuration, this is converted to **dsap isis** on the device.
    - **0xff**—raw IPX 802.3 format. When you apply the configuration, this is converted to **dsap raw-ipx** on the device.
  - **eii-ipx**—Ethernet II IPX format, EtherType 0x8137.
  - **ipx**—Internet Packet Exchange (IPX). This keyword is a shortcut for configuring three separate rules, for **dsap ipx**, **dsap raw-ipx**, and **eii-ipx**. The conversion is made when you apply the configuration to the device.
  - **isis**—Intermediate System to Intermediate System (IS-IS). When you apply the configuration, this keyword is converted to **dsap isis** on the device.
  - **mpls-multicast**— MPLS multicast.
  - **mpls-unicast**—MPLS unicast.
  - *hex_number*—Any EtherType that can be identified by a 16-bit hexadecimal number 0x600 to 0xffff. See RFC 1700, “Assigned Numbers,” at http://www.ietf.org/rfc/rfc1700.txt for a list of EtherTypes.
- **Description**—A explanation of the purpose of the rule, up to 100 characters per line. You can enter multiple lines; each line is added as a remark in the CLI, and the remarks are placed before the rule.
- **More Options** > **Direction**—Whether the rule is for the **In** or **Out** direction. **In** is the default.

When you are finished defining the rule, click **OK** to add the rule to the table.

Step 4 Click **Apply** to save the rule to your configuration.

Configure ICMP Access Rules

By default, you can send ICMP packets to any interface using either IPv4 or IPv6, with these exceptions:

- The ASA does not respond to ICMP echo requests directed to a broadcast address.
- The ASA only responds to ICMP traffic sent to the interface that traffic comes in on; you cannot send ICMP traffic through an interface to a far interface.

To protect the device from attacks, you can use ICMP rules to limit ICMP access to interfaces to particular hosts, networks, or ICMP types. ICMP rules function like access rules, where the rules are ordered, and the first rule that matches a packet defines the action.

If you configure any ICMP rule for an interface, an implicit deny ICMP rule is added to the end of the ICMP rule list, changing the default behavior. Thus, if you want to simply deny a few message types, you must include a permit any rule at the end of the ICMP rule list to allow the remaining message types.

We recommend that you always grant permission for the ICMP unreachable message type (type 3). Denying ICMP unreachable messages disables ICMP path MTU discovery, which can halt IPsec and PPTP traffic. Additionally ICMP packets in IPv6 are used in the IPv6 neighbor discovery process.

Procedure

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

Step 2 Configure ICMP rules:

a) Add a rule (**Add > Rule**, **Add > IPv6 Rule**, or **Add > Insert**), or select a rule and edit it.
b) Select the ICMP type you want to control, or **any** to apply to all types.
c) Select the interface to which the rule applies. You must create separate rules for each interface.
d) Select whether you are permitting or denying access for matching traffic.
e) Select **Any Address** to apply the rule to all traffic. Alternatively, enter the address and mask (for IPv4) or address and prefix length (for IPv6) of the host or network you are trying to control.
f) Click **OK**.

Step 3 (Optional) To set ICMP unreachable message limits, set the following options. Increasing the rate limit, along with enabling the **Decrement time to live for a connection** option in a service policy (on the Configuration > Firewall > Service Policy Rules > Rule Actions > Connection Settings dialog box), is required to allow a trace route through the ASA that shows the ASA as one of the hops.

- Rate Limit—Sets the rate limit of unreachable messages, between 1 and 100 messages per second. The default is 1 message per second.
- Burst Size—Sets the burst rate, between 1 and 10. The system sends this number of replies, but subsequent replies are not sent until the rate limit is reached.

Step 4 Click **Apply**.

Monitoring Access Rules

The Access Rules page includes hit counts for each rule. Mouse over the hit count to see the update time and interval for the count. To reset the hit count, right click the rule and select **Clear Hit Count**, but be aware that this clears the count for all rules applied to the same interface in the same direction.

Evaluating Syslog Messages for Access Rules

Use a syslog event viewer, such as the one in ASDM, to view messages related to access rules.

If you use default logging, you see syslog message 106023 for explicitly denied flows only. Traffic that matches the "implicit deny" entry that ends the rule list is not logged.

If the ASA is attacked, the number of syslog messages for denied packets can be very large. We recommend that you instead enable logging using syslog message 106100, which provides statistics for each rule (including permit rules) and enables you to limit the number of syslog messages produced. Alternatively, you can disable all logging for a given rule.

When you enable logging for message 106100, if a packet matches an ACE, the ASA creates a flow entry to track the number of packets received within a specific interval. The ASA generates a syslog message at the first hit and at the end of each interval, identifying the total number of hits during the interval and the time stamp for the last hit. At the end of each interval, the ASA resets the hit count to 0. If no packets match the ACE during an interval, the ASA deletes the flow entry. When you configure logging for a rule, you can control the interval and even the severity level of the log message, per rule.

A flow is defined by the source and destination IP addresses, protocols, and ports. Because the source port might differ for a new connection between the same two hosts, you might not see the same flow increment because a new flow was created for the connection.

Permitted packets that belong to established connections do not need to be checked against ACLs; only the initial packet is logged and included in the hit count. For connectionless protocols, such as ICMP, all packets are logged, even if they are permitted, and all denied packets are logged.

See the *syslog messages guide* for detailed information about these messages.



Tip When you enable logging for message 106100, if a packet matches an ACE, the ASA creates a flow entry to track the number of packets received within a specific interval. The ASA has a maximum of 32 K logging flows for ACEs. 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 on deny flows only (not on permit flows) because they can indicate an attack. When the limit is reached, the ASA does not create a new deny flow for logging until the existing flows expire, and issues message 106101. You can control the frequency of this message, and the maximum number of deny flows cached, in the advanced settings; see Configure Advanced Options for Access Rules, on page 12.

History for Access Rules

Feature Name	Platform Releases	Description
Interface access rules	7.0(1)	Controlling network access through the ASA using ACLs. We introduced the following screen: Configuration > Firewall > Access Rules.
Global access rules	8.3(1)	Global access rules were introduced. 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	8.4(5), 9.1(2)	In transparent firewall mode, the ASA can now pass IS-IS traffic using an EtherType ACL. We modified the following screen: Configuration > Device Management > Management Access > EtherType Rules.
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.
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. We modified the following screens: Configuration > Firewall > Access Rules Configuration > Remote Access VPN > Network (Client) Access > Group Policies > General > More Options
Extended ACL and object enhancement to filter ICMP traffic by ICMP code	9.0(1)	ICMP traffic can now be permitted/denied based on ICMP code. We introduced or modified the following screens: Configuration > Firewall > Objects > Service Objects/Groups Configuration > Firewall > Access Rule
Transactional Commit Model on Access Group Rule Engine	9.1(5)	When enabled, a rule update is applied after the rule compilation is completed; without affecting the rule matching performance. We introduced the following screen: Configuration > Device Management > Advanced > Rule Engine.

Feature Name	Platform Releases	Description
Configuration session for editing ACLs and objects. Forward referencing of objects and ACLs in access rules.	9.3(2)	You can now edit ACLs and objects in an isolated configuration session. You can also forward reference objects and ACLs, that is, configure rules and access groups for objects or ACLs that do not yet exist.
Access rule support for Stream Control Transmission Protocol (SCTP)	9.5(2)	You can now create access rules using the **sctp** protocol, including port specifications. We modified the add/edit dialog boxes for access rules on the **Configuration** > **Firewall** > **Access Rules** page.
Ethertype rule support for the IEEE 802.2 Logical Link Control packet's Destination Service Access Point address.	9.6(2)	You can now write Ethertype access control rules for the IEEE 802.2 Logical Link Control packet's Destination Service Access Point address. Because of this addition, the **bpdu** keyword no longer matches the intended traffic. Rewrite **bpdu** rules for **dsap 0x42**. We modified the following screen: **Configuration** > **Firewall** > **EtherType Rules**.
Support in routed mode for Ethertype rules on bridge group member interfaces and extended access rules on Bridge Group Virtual Interfaces (BVI).	9.7(1)	You can now create Ethertype ACLs and apply them to bridge group member interfaces in routed mode. You can also apply extended access rules to the Bridge Virtual Interface (BVI) in addition to the member interfaces. We modified the following screens: **Configuration** > **Firewall** > **Access Rules**, **Configuration** > **Firewall** > **Ethertype Rules**.
EtherType access control list changes.	9.9(1)	EtherType access control lists now support Ethernet II IPX (EII IPX). In addition, new keywords are added to the DSAP keyword to support common DSAP values: BPDU (0x42), IPX (0xE0), Raw IPX (0xFF), and ISIS (0xFE). Consequently, existing EtherType access control entries that use the BPDU or ISIS keywords will be converted automatically to use the DSAP specification, and rules for IPX will be converted to 3 rules (DSAP IPX, DSAP Raw IPX, and EII IPX). In addition, packet capture that uses IPX as an EtherType value has been deprecated, because IPX corresponds to 3 separate EtherTypes. We modified the following screens: **Configuration** > **Firewall** > **Ethertype Rules**.
The object group search threshold is now disabled by default.	9.12(1)	If you enabled object group search, the feature was subject to a threshold to help prevent performance degradation. That threshold is now disabled by default. You can enable it by using the **object-group-search threshold** command. We changed the following screen: **Configuration** > **Access Rules** > **Advanced**.

Feature Name	Platform Releases	Description
Forward referencing of ACLs and objects is always enabled. In addition, object group search for access control is now enabled by default.	9.18(1)	You can refer to ACLs or network objects that do not yet exist when configuring access groups or access rules. In addition, object group search is now enabled by default for access control for *new* deployments. Upgrading devices will continue to have this command disabled. If you want to enable it (recommended), you must do so manually. We removed the **forward-reference enable** command, and changed the default for **object-group-search access-control** to enabled.