CHAPTER 32

Configuring Access Rules

This chapter describes how to control network access through the adaptive security appliance using access rules, and it 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 adaptive security appliance 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 34, "Configuring Management Access."

Information About Access Rules

You create an access rule by applying an extended or EtherType access list to an 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:

- IPv4 traffic from a higher security interface to a lower security interface.
- 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:

- IPv4 traffic from a higher security interface to a lower security interface.
- 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 extended access rule (IPv4), an IPv6 access rule (IPv6), or an EtherType rule (non-IPv4/IPv6).

Implicit Deny

Interface-specific access rules do not have an implicit deny at the end, but global rules on inbound traffic do 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 adaptive security appliance except for particular addresses, then you need to deny the particular addresses and then permit all others.

When you have no global access rules in your configuration, the implicit deny rule is applied at the end of interface access rules. When you configure both an interface access rule and a global access rule, the implicit deny (any any) is no longer located at the end of the interface-based access rule. The implicit deny (any any) is enforced at the end of the global access rule. Logically, the entries on the interface-based access rule are processed first, followed by the entries on the global access rule, and then finally the implicit deny (any any) at the end of the global access rule.

For example, when you have an interface-based access rule and a global access rule in your configuration, the following processing logic applies:

1. interface access control rules

2. global access control rules
3. default global access control rule (deny any any)

When only interface-based access rules are configured, the following processing logic applies:

1. interface access control rules
2. default interface access control rule (deny any any)

For EtherType rules, the implicit deny does not affect IPv4 or IPv6 traffic or ARPs; for example, if you allow EtherType 8037 (the EtherType for IPX), the implicit deny at the end of the list does not block any IP traffic that you previously allowed with an access rule (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.

Using Access Rules and EtherType Rules on the Same Interface

You can apply one access rule and one EtherType rule to each direction of an interface.

Inbound and Outbound Rules

The adaptive security appliance supports two types of access lists:

- Inbound—Inbound access lists apply to traffic as it enters an interface.
- 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 adaptive security appliance on an interface or traffic exiting the adaptive security appliance 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 inbound access list can bind an access list to a specific interface or apply a global rule on all interfaces. For more information about global rules, see the "Using Global Access Rules" section on page 32-4.

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 32-1.) The outbound access list prevents any other hosts from reaching the outside network.

Figure 32-1 Outbound Access List

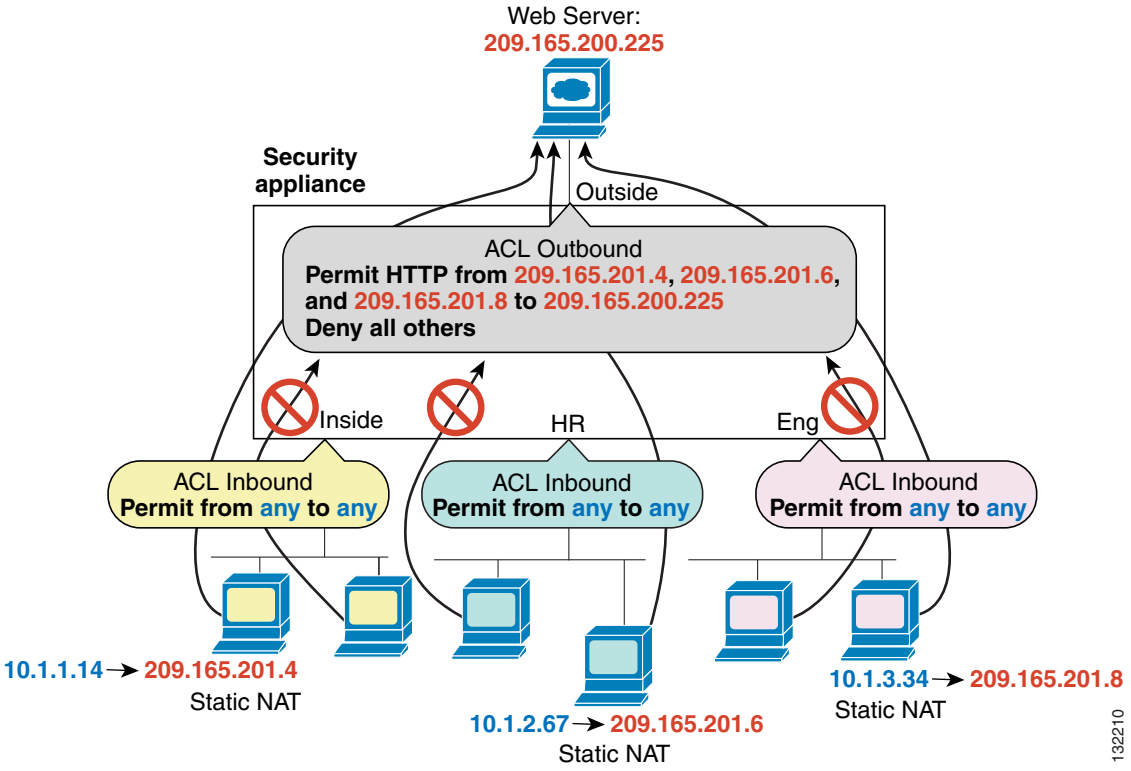

See the following commands for this example:

```
hostname(config)# access-list OUTSIDE extended permit tcp host 209.165.201.4
host 209.165.200.225 eq www
hostname(config)# access-list OUTSIDE extended permit tcp host 209.165.201.6
host 209.165.200.225 eq www
hostname(config)# access-list OUTSIDE extended permit tcp host 209.165.201.8
host 209.165.200.225 eq www
hostname(config)# access-group OUTSIDE out interface outside
```

Using Global Access Rules

Global access rules allow you to apply a global rule to ingress traffic without the need to specify an interface to which the rule must be applied. Using global access rules provides the following benefits:

- When migrating to the adaptive security appliance from a competitor appliance, you can maintain a global access rule policy instead of needing to apply an interface-specific policy on each interface.
- Global access control policies are not replicated on each interface, so they save memory space.
- Global access rules provides flexibility in defining a security policy. You do not need to specify which interface a packet comes in on, as long as it matches the source and destination IP addresses.
- Global access rules use the same mtrie and stride tree as interface-specific access rules, so scalability and performance for global rules are the same as for interface-specific rules.

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.

For information about implicit deny in global access rules, see the "Implicit Deny" section on page 32-2.

Information About Extended Access Rules

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

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 adaptive security appliance allows all returning traffic for established, bidirectional connections.

For connectionless protocols such as ICMP, however, the adaptive security appliance 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. To control ping, specify **echo-reply** (**0**) (adaptive security appliance to host) or **echo** (**8**) (host to adaptive security appliance).

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. This feature is especially useful in multiple context mode, which does not allow dynamic routing, for example.

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

Table 32-1 lists common traffic types that you can allow through the transparent firewall.

Table 32-1 Transparent Firewall Special Traffic

Traffic Type	Protocol or Port	Notes
DHCP	UDP ports 67 and 68	If you enable the DHCP server, then the adaptive security appliance 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 adaptive security appliance. Access control rules for to-the-box management traffic (defined by such commands as **http**, **ssh**, or **telnet**) have higher precedence than an management access rule applied with the **control-plane** option. 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

- An EtherType rule controls any EtherType identified by a 16-bit hexadecimal number.
- EtherType rules support Ethernet V2 frames.
- 802.3-formatted frames are not handled by the rule because they use a length field as opposed to a type field.
- BPDUs, which are permitted by default, are the only exception: they are SNAP-encapsulated, and the adaptive security appliance is designed to specifically handle BPDUs.
- The adaptive security appliance receives trunk port (Cisco proprietary) BPDUs. Trunk BPDUs have VLAN information inside the payload, so the adaptive security appliance modifies the payload with the outgoing VLAN if you allow BPDUs.

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 adaptive security appliance by configuring both MPLS routers connected to the adaptive security appliance to use the IP address on the adaptive security appliance 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 adaptive security appliance.

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

Or

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

Licensing Requirements for Access Rules

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Prerequisites

Before you can create an access rule, create the access list. See Chapter 13, "Adding an Extended Access List," and Chapter 14, "Adding an EtherType Access List," for more information.

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

Per-User Access List Guidelines

- If there is no per-user access list associated with a packet, the interface access rule is applied.
- The per-user access list 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 access list, syslog message 109025 is logged. If traffic is permitted, no syslog message is generated. The **log** option in the per-user access list has no effect.

Additional Guidelines and Limitations

To access the adaptive security appliance interface for management access, you do not need an access list allowing the host IP address. You only need to configure management access by following the instructions in Chapter 34, "Configuring Management Access."

Default Settings

See the "Implicit Permits" section on page 32-2.

Configuring Access Rules

To apply an access rule, perform the following steps.

Command	Purpose
access-group *access_list* {{**in** | **out**} **interface** *interface_name* [**per-user-override** | **control-plane**] | **global**} **Example:** `hostname(config)# access-group acl_out in interface outside`	Binds an access list to an interface or applies it globally. **Note** The **access-group** command cannot reference empty access lists or access lists that contain only a remark. The **in** keyword applies the access list to the traffic on the specified interface. The **out** keyword applies the access list to the outbound traffic. The **per-user-override** keyword (for inbound access lists only) allows dynamic user access lists that are downloaded for user authorization to override the access list assigned to the interface. For example, if the interface access list denies all traffic from 10.0.0.0, but the dynamic access list permits all traffic from 10.0.0.0, then the dynamic access list overrides the interface access list for that user. See the "Configuring RADIUS Authorization" section on page 35-10 for more information about per-user access lists. See also the "Per-User Access List Guidelines" section on page 32-7. The **control-plane** keyword specifies if the rule is for to-the-box traffic. The **global** keyword applies the access list to the inbound direction of all interfaces.

Examples

The following example shows how to use the **access-group** command:

```
hostname(config)# access-list acl_out permit tcp any host 209.165.201.3 eq 80
hostname(config)# access-group acl_out in interface outside
```

The **access-list** command lets any host access the global address using port 80. The **access-group** command specifies that the **access-list** command applies to traffic entering the outside interface.

Monitoring Access Rules

To monitor network access, perform one of the following tasks:

Command	Purpose
show running-config access-group	Displays the current access list bound to the interfaces.

Configuration Examples for Permitting or Denying Network Access

This section includes typical configuration examples for permitting or denying network access.

The following example illustrates the commands required to enable access to an inside web server with the IP address 209.165.201.12. (This IP address is the real address, not the visible on the outside interface after NAT.)

```
hostname(config)# access-list ACL_OUT extended permit tcp any host 209.165.201.12 eq www
hostname(config)# access-group ACL_OUT in interface outside
```

The following example allows all hosts to communicate between the **inside** and **hr** networks but only specific hosts to access the outside network:

```
hostname(config)# access-list ANY extended permit ip any any
hostname(config)# access-list OUT extended permit ip host 209.168.200.3 any
hostname(config)# access-list OUT extended permit ip host 209.168.200.4 any

hostname(config)# access-group ANY in interface inside
hostname(config)# access-group ANY in interface hr
hostname(config)# access-group OUT out interface outside
```

For example, the following sample access list allows common EtherTypes originating on the inside interface:

```
hostname(config)# access-list ETHER ethertype permit ipx
hostname(config)# access-list ETHER ethertype permit mpls-unicast
hostname(config)# access-group ETHER in interface inside
```

The following example allows some EtherTypes through the adaptive security appliance, but it denies all others:

```
hostname(config)# access-list ETHER ethertype permit 0x1234
hostname(config)# access-list ETHER ethertype permit mpls-unicast
hostname(config)# access-group ETHER in interface inside
hostname(config)# access-group ETHER in interface outside
```

The following example denies traffic with EtherType 0x1256 but allows all others on both interfaces:

```
hostname(config)# access-list nonIP ethertype deny 1256
hostname(config)# access-list nonIP ethertype permit any
hostname(config)# access-group ETHER in interface inside
hostname(config)# access-group ETHER in interface outside
```

The following example uses object groups to permit specific traffic on the inside interface:

```
!
hostname (config)# object-group service myaclog
hostname (config-service)# service-object tcp source range 2000 3000
hostname (config-service)# service-object tcp source range 3000 3010 destinatio$
hostname (config-service)# service-object ipsec
hostname (config-service)# service-object udp destination range 1002 1006
hostname (config-service)# service-object icmp echo

hostname(config)# access-list outsideacl extended permit object-group myaclog interface
inside any
```

Feature History for Access Rules

Table 32-2 lists each feature change and the platform release in which it was implemented.

Table 32-2 Feature History for Access Rules

Feature Name	Platform Releases	Feature Information
Interface access rules.	7.0(1)	Controlling network access through the security appliance using access lists. The following command was introduced: **access-group**.
Global access rules.	8.3(1)	Global access rules were introduced. The following command was modified: **access-group**.