



Configuring Reflexive Access Lists

- [Restrictions on Using Reflexive Access Lists, on page 1](#)
- [Information About Reflexive Access Lists, on page 1](#)
- [How to Configure Reflexive Access Lists, on page 7](#)
- [Configuration Examples for Reflexive Access List, on page 10](#)
- [Feature History for Reflexive Access Lists, on page 12](#)

Restrictions on Using Reflexive Access Lists

- Reflexive access lists do not work with some applications that use port numbers that change during a session. For example, if the port numbers for a return packet are different from the originating packet, reflexive access list denies the return packet. Even if the return packet is actually a part of the same session. The TCP application of FTP is an example of an application with changing port numbers. If you start an FTP request from within your network, reflexive access lists will not complete the request. Instead, use passive FTP when sending requests from within your network.
- Wireshark functionality will not work for packets that are filtered through reflexive access lists.
- Netflow-based features (such as NetFlow, Encrypted Traffic Analytics [ETA], Wired Application Visibility and Control [WDAVC], Security Group Tag Caching [SGT-C]) and reflexive access lists can't be configured on the same interface.
- Reflexive access list configuration supports only 5-tuple with port range.
- Only Layer3 interfaces (L3 interfaces, L3 subinterfaces, SVIs, L3 port channel, and port-channel subinterfaces) support reflexive access list configuration.

Information About Reflexive Access Lists

Reflexive access lists allow filtering of IP packets based on upper-layer session information. You can use reflexive access lists to permit IP traffic for sessions originating within your network. You can also deny IP traffic for sessions originating outside your network. Reflexive access lists accomplish this by using reflexive filtering, a kind of session filtering.

Reflexive access lists can be defined with extended, named IP access lists only. You cannot define reflexive access lists with numbered or standard, named IP access lists or with other protocol access lists.

You can use reflexive access lists along with other standard access lists and static, extended access lists.



Note Reflexive access lists supports both ingress NetFlow and egress NetFlow. The default size of egress NetFlow with custom SDM template is zero, and with custom SDM template the NetFlow table size can be set to zero for both directions. In such cases, reflexive ACL does not work as NetFlow cannot be installed in the hardware flow table.

Benefits of Reflexive Access Lists

Reflexive access lists are an important part of securing your network against network hackers, and can be included in a firewall defense. Reflexive access lists provide a level of security against spoofing and certain denial-of-service attacks. Reflexive access lists are simple to use, and, compared to basic access lists, provide greater control over which packets enter your network.

Overview of Reflexive Access Lists

Reflexive access lists are similar in many ways to other access lists. Reflexive access lists contain condition statements that define criteria for permitting IP packets. These entries are evaluated in the order in which they are entered in the list. When a match occurs, no more entries are evaluated.

However, reflexive access lists have significant differences from other types of access lists. Reflexive access lists contain only temporary entries. These entries are automatically created when a new IP session begins. The entries are removed when the session ends. Reflexive access lists aren't themselves applied directly to an interface. They are nested within an extended, named IP access list that is applied to the interface. For more information about this, see [How to Configure Reflexive Access Lists, on page 7](#). Also, reflexive access lists don't have the usual implicit `deny all traffic` statement at the end of the list because they are nested lists.

Implementing Session Filtering with Reflexive Access Lists

The following section provides information on how reflexive access lists are better at implementing session filtering.

Session Filtering with Basic Access Lists

With standard and extended access lists, you can implement a basic version of session filtering by using the **established** keyword with the **permit** command. The **established** keyword filters TCP packets based on whether the ACK or RST bits are set. Set ACK or RST bits indicate that the packet isn't the first in the session. Therefore, the packet belongs to an established session. This filter criterion is a part of an access list applied permanently to an interface.

This method of using the **established** keyword is available only for the TCP upper-layer protocol. For the other upper-layer protocols, such as UDP, ICMP, you have to either permit all the incoming traffic or define all the possible permissible source/destination host/port address pairs for each protocol. This is an unmanageable task and might exhaust the NVRAM space.

Session Filtering with Reflexive Access Lists

Reflexive access lists provide a truer form of session filtering. This filtering is much harder to spoof because more filter criteria must be matched before a packet is permitted through. For example, source and destination addresses and port numbers are checked, and not just the ACK and RST bits. Also, session filtering uses temporary filters that are removed after a session is over. This limits a hacker's attack opportunity to a smaller time window.

Location at which to Configure Reflexive Access Lists

Configure reflexive access lists on border devices and devices that pass traffic between an internal and external network.



Note In this chapter, the phrases 'within your network' and 'internal network' refer to a controlled and secured network, for example, the intranet of your organization. They also refer to a part of the internal network of your organization that has higher security requirements than another part. 'Outside your network' and 'external network' refer to a network that is uncontrolled and unsecured, for example the internet. They also refer to a part of your organization's network that isn't highly secured.

How Reflexive Access Lists Work

A reflexive access list is triggered when a new IP upper-layer session (such as TCP or UDP) is initiated from inside your network, with a packet traveling to the external network. When triggered, the reflexive access list generates a new, temporary entry. This entry permits traffic to enter your network if the traffic is a part of the session. The entry won't permit traffic to enter your network if the traffic isn't a part of the session.

For example, when the first outbound packet of a TCP session is forwarded outside of your network, a new temporary reflexive access list entry is created. This entry is added to the reflexive access list, that applies to inbound traffic.

The number of configurable reflexive access list groups is limited to 100. Each group can have 100 ACL entries. Theoretically, while the number of ACLs can be 10,000, the actual number depends on the TCAM size supported by your device. The sequence number of the Reflexive ACEs can't exceed 65530. When the limit is exceeded, the forwarding manager logs an error message on the console indicating the same.

The temporary reflexive access list entries are learned using the NetFlow hash table in the hardware. If the flow table is full when a new entry is detected, the traffic for the failed flow is dropped.

Dynamic reflexive entries aren't synced to the standby. When Stateful Switchover takes place in a StackWise Virtual system, notifications are sent to clear the dynamic entries from the hardware when the system role changes from standby to active. The temporary reflexive access list entries are relearned.

Temporary Access List Entry Characteristics

- The entry is always a **permit** entry.
- The entry specifies the same protocol, such as TCP, as the original outbound packet.
- The entry specifies the same source and destination addresses as the original outbound TCP packet, except that the addresses are swapped.

- The entry specifies the same source and destination port numbers as the original outbound TCP packet, except that the port numbers are swapped.

This entry characteristic applies only for TCP and UDP packets. Other protocols, such as ICMP and IGMP, don't have port numbers, and other criteria are specified. For example, for ICMP, type numbers are used.

- Inbound TCP traffic is evaluated against the entry until the entry expires. If an inbound TCP packet matches the entry, the inbound packet is forwarded into your network.
- The entry expires after the last packet of the session passes through the interface.
- If no packets belonging to the session are detected for the timeout period, the entry expires.
- The per entry timeout minimum value should align with the global timeout value. Per entry timeout value that is lower than the global timeout value won't be honored.

Temporary reflexive access list entries are removed at the end of the session. For TCP sessions, the entry is removed 5 seconds after two set FIN bits are detected, or immediately after a TCP packet is matched with the RST bit set. Two set FIN bits in a session indicate that the session is about to end; the 5-second window allows the session to close gracefully. A set RST bit indicates an abrupt session close. The temporary entry is removed when no packets of the session are detected for the timeout period.

For UDP and other protocols, the end of the session is determined differently than for a TCP session. Because other protocols are considered to be sessionless services, no session-tracking information is embedded in packets. Therefore, a session is considered to have ended when no packets of the session are detected for the timeout period.

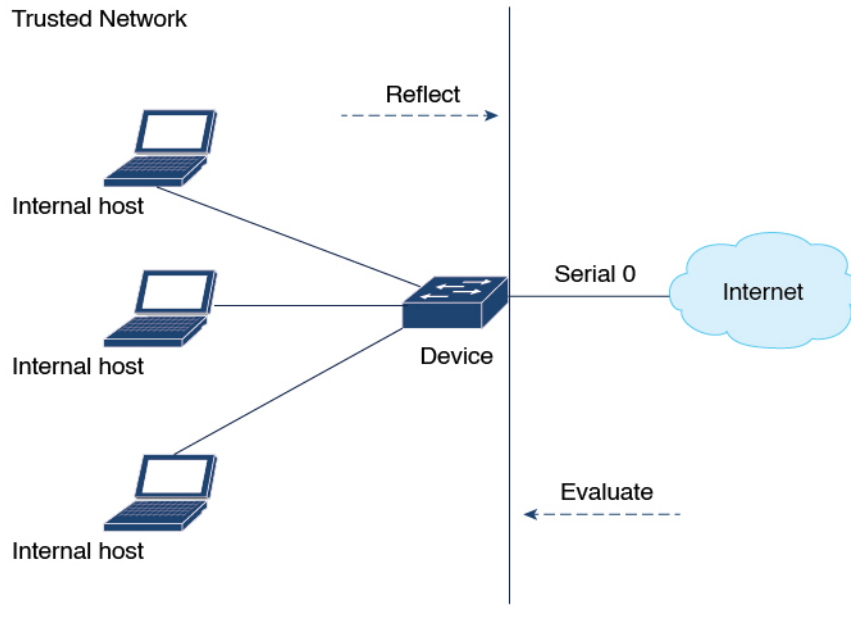
Choosing an Interface Internal or External

Before you configure reflexive access lists, you must decide whether to configure reflexive access lists on an internal or external interface. You should also be sure that you have a basic understanding of the IP protocol and of access lists. Specifically, you should know how to configure extended named IP access lists. To learn about configuring IP extended access lists, refer to the “Configuring IP Services” chapter of the *Cisco IOS IP Configuration Guide*.

Reflexive access lists are most commonly used with one of two basic network topologies. Determining which of these topologies is most like your own helps you decide whether to use reflexive access lists with an internal interface or with an external interface. An internal interface is the interface connecting to an internal network. An external interface is the interface connecting to an external network.

In the first topology, reflexive access lists are configured for the external interface. This prevents IP traffic from entering the device and the internal network, unless the traffic is part of a session already established from within the internal network.


Figure 1: Reflexive Access List on the External Interface



357951

In the second topology, reflexive access lists are configured for the internal interface. This allows external traffic to access the services in the Demilitarized Zone (DMZ), such as DNS services. But, prevents IP traffic from entering your internal network - unless the traffic is part of a session already established from within the internal network.

Figure 2: Reflexive Access List for the Internal Interface



357950

External Interface Configuration Task List

To configure reflexive access lists for an external interface, perform the following tasks:

1. Define the reflexive access lists in an outbound IP extended named access list.
2. Nest the reflexive access lists in an inbound IP extended named access list.
3. Set a global timeout value.

These tasks are described in the section [Defining A Reflexive Access List, on page 7](#).



Note The outbound reflexive access list evaluates traffic traveling out of your network. If the defined reflexive access list is matched, temporary entries are created in the nested, inbound reflexive access list. These temporary entries are then applied to the traffic traveling into your network.

Internal Interface Configuration Task List

To configure reflexive access lists for an internal interface, perform the following tasks:

1. Define the reflexive access lists in an inbound IP extended named access list.
2. Nest the reflexive access lists in an outbound IP extended named access list.
3. Set a global timeout value

These tasks are described in the [Defining A Reflexive Access List, on page 7](#) section.



Note The inbound reflexive access list is used to evaluate traffic traveling out of your network. If the defined reflexive access list is matched, temporary entries are created in the nested, outbound reflexive access list. These temporary entries are then applied to the traffic traveling into your network.

Mixing Reflexive Access List Statements with Other Permit and Deny Entries

The extended IP access list that contains the reflexive access list **permit** statement can also contain other normal **permit** and **deny** statements. However, as with all access lists, the order of entries is important.

When an outbound IP packet reaches an external interface configured with a reflexive access list, the packet is evaluated sequentially by each entry in the outbound access list until a match occurs.

If the packet matches an entry prior to the reflexive **permit** entry, the packet won't be evaluated by the reflexive **permit** entry. No temporary entry is created for the reflexive access list.

The outbound packet is evaluated by the reflexive **permit** entry only if no other match occurs first. Then, if the packet matches the protocol specified in the reflexive **permit** entry, the packet is forwarded out of the interface. A corresponding temporary entry is created in the inbound reflexive access list, unless the corresponding entry already exists. If the entry already exists it would indicate that the outbound packet belongs to a session in progress. The temporary entry specifies criteria that permit inbound traffic for the same session only.

How to Configure Reflexive Access Lists

The following sections describe the procedures you can perform to define, nest and clear reflexive access lists and to set the global timeout value.

Defining A Reflexive Access List

To define a reflexive access list, use an entry in an extended named IP access list. This entry must use the **reflect** keyword.

- If you are configuring reflexive access lists for an external interface, apply the extended named IP access list to outbound traffic.
- If you are configuring reflexive access lists for an internal interface, apply the extended named IP access list to inbound traffic.
- If you specify an extended named IP access list, ensure that you apply the list to the interface.

Procedure

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | enable Example: Device> enable | Enables privileged EXEC mode. Enter your password, if prompted. |
| Step 2 | configure terminal Example: Device# configure terminal | Enters global configuration mode. |
| Step 3 | ip access-list extended name Example: Device(config)# ip access-list extended | Enters the access-list configuration mode. Specifies the outbound access list for an external interface. Or Specifies the inbound access list for an internal interface. |
| Step 4 | permit protocol-name any any reflect name [timeout seconds] Example: Device(config-ext-nacl)# permit tcp any any reflect tcptraffic [timeout 20] | Defines the reflexive access list using the reflexive permit entry. Repeat this step for each IP upper-layer protocol, for example, you can define reflexive filtering for TCP sessions and for UDP sessions. You can use the same <i>name</i> for multiple protocols. |

| | Command or Action | Purpose |
|---------------|---|--|
| | | Note You can create more than one reflexive access list per ACL. You can have several reflexive lists that can be tied in to any number of items in the ACL. The ACL items can be common to one input interface or many interfaces. They can be evaluated on different output interfaces. |
| Step 5 | exit Example: Device(config-ext-nacl) # exit | Exits access-list configuration mode and enters global configuration mode. |
| Step 6 | interface <i>type number</i> Example: Device(config) # interface GigabitEthernet 1 | Configures an interface and enters interface configuration mode. |
| Step 7 | ip access-group <i>name out</i> Example: Device(config-if) # ip access-group outboundfilters out | Applies the extended access list to the outbound traffic of the interface. |
| Step 8 | ip access-group <i>name in</i> Example: Device(config-if) # ip access-group inboundfilters in | Applies the extended access list to the inbound traffic of the interface. |

Nesting Reflexive Access Lists

After you define a reflexive access list in one IP extended access list, you must nest the reflexive access list within a different extended named IP access list:

- When you configure reflexive access lists for an external interface, nest the reflexive access list within an extended named IP access list applied to inbound traffic.
- When you configure reflexive access lists for an internal interface, nest the reflexive access list within an extended named IP access list applied to outbound traffic.

After you nest a reflexive access list, packets heading into your internal network can be evaluated against any reflexive access list temporary entries, along with the other entries in the extended named IP access list.

Again, the order of entries is important. Normally, when a packet is evaluated against entries in an access list, the entries are evaluated in sequential order. When a match occurs, no more entries are evaluated. With a reflexive access list nested in an extended access list, the extended access list entries are evaluated sequentially up to the nested entry. Then the reflexive access list entries are evaluated sequentially. And then the remaining entries in the extended access list are evaluated sequentially. As usual, after a packet matches any of these entries, no more entries are evaluated.

When you specify an extended named IP access list, ensure that you apply the list to the interface.

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | enable Example: Device> enable | Enables privileged EXEC mode. Enter your password, if prompted. |
| Step 2 | configure terminal Example: Device# configure terminal | Enters global configuration mode. |
| Step 3 | ip access-list extended <i>name</i> Example: Device(config)# ip access-list extended outboundfilters | Enters the access-list configuration mode. Specifies the outbound access list for an external interface. Or Specifies the inbound access list for an internal interface. |
| Step 4 | evaluate <i>name</i> Example: Device(config-ext-nacl)# evaluate tcptraffic | Adds an entry that points to the reflexive access list. An entry is added for each reflexive access list with the name <i>name</i> previously defined. |
| Step 5 | exit Example: Device(config-ext-nacl)# exit | Exits access-list configuration mode and enters global configuration mode. |
| Step 6 | interface <i>type number</i> Example: Device(config)# interface type number | Configures an interface and enters interface configuration mode. |
| Step 7 | ip access-group <i>name out</i> Example: Device(config-if)# ip access-group outboundfilters out | Applies the extended access list to the outbound traffic of the interface. |
| Step 8 | ip access-group <i>namein</i> Example: Device(config-if)# ip access-group inboundfilters in | Applies the extended access list to the inbound traffic of the interface. |

Setting a Global Timeout Value

Reflexive access list entries expire when no packets are detected for a certain length of time in the session. This length of time is called the timeout period. You can specify the timeout period for a particular reflexive access list when you define the reflexive access list. If you do not specify the timeout period for a given reflexive access list, the list uses the global timeout value instead.

The global timeout value is 300 seconds by default. You can change the global timeout value any time.

To change the global timeout value, use the following command in global configuration mode.

| Command | Purpose |
|--|--|
| Device (config)# ip reflexive-list timeout <i>seconds</i> | Changes the global timeout value for temporary reflexive access list entries. Use a positive integer from 30 to 2,147,483. |

Clearing a Reflexive Access List

To clear a reflexive access list, perform the following procedure.

Procedure

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | enable Example: Device> enable | Enables privileged EXEC mode. Enter your password, if prompted. |
| Step 2 | clear ip reflexive-list { * <i>reflexive-list name</i> } Example: Device# clear ip reflexive-list reflexive1 | Deletes the reflexive access list entries listed in the reflexive access list titled <i>reflexive1</i> . The * keyword deletes the access list entries in all the reflexive access lists. |

Configuration Examples for Reflexive Access List

The following sections provide configuration examples for reflexive access lists.

Example: External Interface Configuration

This example shows configuration of reflexive access lists for an external interface.

This configuration example permits both inbound and outbound TCP traffic at interface GigabitEthernet 1, but only if the first packet in a given session originated from inside your network. Interface GigabitEthernet 1 connects to the internet.

This command defines the interface where session-filtering configuration is to be applied:

```
interface GigabitEthernet 1
  description Access to the Internet via this interface
```

This command applies access lists to an interface, for inbound traffic and outbound traffic.

```
ip access-group inboundfilters in
ip access-group outboundfilters out
```

This command defines an outbound access list. This is the access list that evaluates all outbound traffic on interface GigabitEthernet 1.

```
ip access-list extended outboundfilters
```

This command defines a reflexive access list called `tcptraffic`. This entry permits all outbound TCP traffic and creates a new access list named `tcptraffic`. Also, when an outbound TCP packet is the first in a new session, a corresponding temporary entry is automatically created in the reflexive access list `tcptraffic`.

```
permit tcp any any reflect tcptraffic
```

This command defines an inbound access list. This is the access list that evaluates all inbound traffic on interface GigabitEthernet 1.

```
ip access-list extended inboundfilters
```

This command defines inbound access list entries. This example shows that EIGRP is permitted on the interface. Also, no ICMP traffic is permitted. The last entry points to the reflexive access list. If a packet doesn't match the first two entries, the packet is evaluated against all the entries in the reflexive access list called `tcptraffic`.

```
permit eigrp any any
deny icmp any any
evaluate tcptraffic
```

This command defines the global idle timeout value for all reflexive access lists. In this example, when the reflexive access list `tcptraffic` was defined, no timeout was specified, so `tcptraffic` uses the global timeout. Therefore, if for 120 seconds there's no TCP traffic that is part of an established session, the corresponding reflexive access list entry will be removed.

```
ip reflexive-list timeout 120
```

The example configuration looks as follows:

```
interface GigabitEthernet 1
  description Access to the Internet via this interface
  ip access-group inboundfilters in
  ip access-group outboundfilters out
  !
  ip reflexive-list timeout 120
  !
  ip access-list extended outboundfilters
  permit tcp any any reflect tcptraffic
  !
  ip access-list extended inboundfilters
  permit eigrp any any
  deny icmp any any
  evaluate tcptraffic
```

In this configuration, before any TCP sessions are initiated, the **show access-list** command displays the following:

```
Extended IP access list inboundfilters
```

```

permit eigrp any any
deny icmp any any
evaluate tcptraffic
Extended IP access list outboundfilters
permit tcp any any reflect tcptraffic

```

The reflexive access list doesn't appear in this output. This is because before any TCP sessions are initiated, no traffic has triggered the reflexive access list, and the list is empty, that is, it has no entries. When empty, reflexive access lists do not show up in the **show access-list** output.

After a Telnet connection is initiated from within your network to a destination outside of your network, the **show access-list** command displays the following:

```

Extended IP access list inboundfilters
permit eigrp any any
deny icmp any any
evaluate tcptraffic
Extended IP access list outboundfilters
permit tcp any any reflect tcptraffic
Reflexive IP access list tcptraffic
permit tcp host 172.19.99.67 eq telnet host 192.168.60.185 eq 11005 (5 matches) (time
left 115 seconds)

```

The reflexive access list `tcptraffic` now appears, and displays the temporary entry generated when the Telnet session is initiated with an outbound packet.

Example: Internal Interface Configuration

The following is an sample configuration for reflexive access lists configured for an internal interface.

```

interface GigabitEthernet 0
description Access from the I-net to our Internal Network via this interface
ip access-group inboundfilters in
ip access-group outboundfilters out
!
ip reflexive-list timeout 120
!
ip access-list extended outboundfilters
permit eigrp any any
deny icmp any any
evaluate tcptraffic
!
ip access-list extended inboundfilters
permit tcp any any reflect tcptraffic

```

Feature History for Reflexive Access Lists

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

| Release | Feature | Feature Information |
|-----------------------------|------------------------|--|
| Cisco IOS XE Dublin 17.10.1 | Reflexive Access Lists | Reflexive access lists allow IP packets to be filtered based on upper-layer session information. |

Use the Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to [Cisco Feature Navigator](#).

