

Configuring IP Session Filtering (Reflexive Access Lists)

This chapter describes how to configure reflexive access lists on your router. Reflexive access lists provide the ability to filter network traffic at a router, based on IP upper-layer protocol “session” information.

For a complete description of reflexive access list commands, refer to the “Reflexive Access List Commands” chapter of the *Security Command Reference*. To locate documentation of other commands that appear in this chapter, use the command reference master index or search online.

In This Chapter

This chapter has the following sections:

- About Reflexive Access Lists
- Prework: Before You Configure Reflexive Access Lists
- Configure Reflexive Access Lists
- Reflexive Access Lists Configuration Examples

About Reflexive Access Lists

Reflexive access lists allow IP packets to be filtered based on upper-layer session information. You can use reflexive access lists to permit IP traffic for sessions originating from within your network but to deny IP traffic for sessions originating from outside your network. This is accomplished by 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 in conjunction with other standard access lists and static extended access lists.

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.

What Is a Reflexive Access List?

Reflexive access lists are similar in many ways to other access lists. Reflexive access lists contain condition statements (entries) that define criteria for permitting IP packets. These entries are evaluated in order, and 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 (for example, with an outbound packet), and the entries are removed when the session ends. Reflexive access lists are not themselves applied directly to an interface, but are “nested” within an extended named IP access list that is applied to the interface. (For more information about this, see the section “Configure Reflexive Access Lists” later in this chapter.) Also, reflexive access lists do not have the usual implicit “deny all traffic” statement at the end of the list, because of the nesting.

How Reflexive Access Lists Implement Session Filtering

This section compares session filtering with basic access lists to session filtering with reflexive access lists.

With Basic Access Lists

With basic standard and static extended access lists, you can approximate 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 is not the first in the session, and therefore, that the packet belongs to an established session.) This filter criterion would be part of an access list applied permanently to an interface.

With Reflexive Access Lists

Reflexive access lists, however, provide a truer form of session filtering, which 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, not just ACK and RST bits.) Also, session filtering uses temporary filters which are removed when a session is over. This limits the hacker’s attack opportunity to a smaller time window.

Moreover, the previous method of using the **established** keyword was available only for the TCP upper-layer protocol. So, for the other upper-layer protocols (such as UDP, ICMP, and so forth), you would have to either permit all incoming traffic or define all possible permissible source/destination host/port address pairs for each protocol. (Besides being an unmanageable task, this could exhaust NVRAM space.)

Where to Configure Reflexive Access Lists

Configure reflexive access lists on border routers—routers that pass traffic between an internal and external network. Often, these are firewall routers.

Note In this chapter, the words “within your network” and “internal network” refer to a network that is controlled (secured), such as your organization’s intranet, or to a part of your organization’s internal network that has higher security requirements than another part. “Outside your network” and “external network” refer to a network that is uncontrolled (unsecured) such as the Internet or to a part of your organization’s network that is not as 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 will permit traffic to enter your network if the traffic is part of the session, but will not permit traffic to enter your network if the traffic is not part of the session.

For example, if an outbound TCP packet is forwarded to outside of your network, and this packet is the first packet of a TCP session, then a new, temporary reflexive access list entry will be created. This entry is added to the reflexive access list, which applies to inbound traffic. The temporary entry has characteristics as described next.

Temporary Access List Entry Characteristics

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

(This entry characteristic applies only for TCP and UDP packets. Other protocols, such as ICMP and IGMP, do not have port numbers, and other criteria are specified. For example, for ICMP, type numbers are used instead.)

- Inbound TCP traffic will be evaluated against the entry, until the entry expires. If an inbound TCP packet matches the entry, the inbound packet will be forwarded into your network.
- The entry will expire (be removed) after the last packet of the session passes through the interface.
- If no packets belonging to the session are detected for a configurable length of time (the timeout period), the entry will expire.

When the Session Ends

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 matching a TCP packet 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.) Or, the temporary entry is removed after no packets of the session have been detected for a configurable length of time (the timeout period).

For UDP and other protocols, the end of the session is determined differently than for TCP. Because other protocols are considered to be connectionless (sessionless) services, there is no session tracking information embedded in packets. Therefore, the end of a session is considered to be when no packets of the session have been detected for a configurable length of time (the timeout period).

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, the return packet will be denied, even if the packet is actually part of the same session.

The TCP application of FTP is an example of an application with changing port numbers. With reflexive access lists, if you start an FTP request from within your network, the request will not complete. Instead, you must use Passive FTP when originating requests from within your network.

Prework: Before You Configure Reflexive Access Lists

Before you configure reflexive access lists, you must decide whether to configure reflexive access lists on an internal or external interface, as described in the next section, “Choose an Interface: Internal or External.”


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 *Network Protocols Configuration Guide, Part 1*.

Choose an Interface: Internal or External

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


The first topology is shown in Figure 7. In this simple topology, reflexive access lists are configured for the *external* interface Serial 1. This prevents IP traffic from entering the router and the internal network, unless the traffic is part of a session already established from within the internal network.

Figure 7 Simple Topology—Reflexive Access Lists Configured at the External Interface



The second topology is shown in Figure 8. In this topology, reflexive access lists are configured for the *internal* interface Ethernet 0. 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 8 DMZ Topology—Reflexive Access Lists Configured at the Internal Interface



Use these two example topologies to help you decide whether to configure reflexive access lists for an internal or external interface.

Configure Reflexive Access Lists

In the previous section, “Prework: Before You Configure Reflexive Access Lists,” you decided whether to configure reflexive access lists for an internal or external interface.

Now, complete the tasks in one of the following configuration task lists.

External Interface Configuration Task List

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

- 1 Define the Reflexive Access List(s) in an *outbound* IP extended named access list.
- 2 Nest the Reflexive Access List(s) in an *inbound* IP extended named access list.
- 3 Set a Global Timeout Value (Optional).

These tasks are described in the sections following the internal interface configuration task list.

Note The defined (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 will then be applied to traffic traveling into your network.

Internal Interface Configuration Task List

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

- 1 Define the Reflexive Access List(s) in an *inbound* IP extended named access list.
- 2 Nest the Reflexive Access List(s) in an *outbound* IP extended named access list.
- 3 Set a Global Timeout Value (Optional).

These tasks are described in the next sections.

Note The defined (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 will then be applied to traffic traveling into your network.

Define the Reflexive Access List(s)

To define a reflexive access list, you 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, the extended named IP access list should be one that is applied to outbound traffic.
- If you are configuring reflexive access lists for an internal interface, the extended named IP access list should be one that is applied to inbound traffic.

To define reflexive access lists, perform the following tasks, starting in global configuration mode:

Task	Command
External interface: Specify the outbound access list.	ip access-list extended <i>name</i>
or	
Internal interface: Specify the inbound access list. (Performing this task also causes you to enter the access-list configuration mode).	
Define 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 also for UDP sessions. You can use the same <i>name</i> for multiple protocols. For additional guidelines for this task, see the following section, “Mixing Reflexive Access List Statements with Other Permit and Deny Entries.”	permit <i>protocol any any reflect</i> <i>name [timeout seconds]</i>

If the extended named IP access list you just specified has never been applied to the interface, you must also complete this next task:

Apply the extended named IP access list to the interface, in interface configuration mode:

Task	Command
External interface: Apply the extended access list to the interface's outbound traffic.	ip access-group <i>name</i> out
or	or
Internal interface: Apply the extended access list to the interface's inbound traffic.	ip access-group <i>name</i> in

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 (entries). However, as with all access lists, the order of entries is important, as explained in the next few paragraphs.

If you configure reflexive access lists for an external interface, when an outbound IP packet reaches the interface, the packet will be 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 will not be evaluated by the reflexive **permit** entry, and no temporary entry will be created for the reflexive access list (reflexive filtering will not be triggered).

The outbound packet will be 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 and a corresponding temporary entry is created in the inbound reflexive access list (unless the corresponding entry already exists, indicating the outbound packet belongs to a session in progress). The temporary entry specifies criteria that permits inbound traffic only for the same session.

Nest the Reflexive Access List(s)

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.

- If you are configuring reflexive access lists for an external interface, nest the reflexive access list within an extended named IP access list applied to inbound traffic.
- If you are configuring 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.

To nest reflexive access lists, perform the following tasks, starting in global configuration mode:

Task	Command
External interface: Specify the inbound access list. or Internal interface: Specify the outbound access list. (Performing this task also causes you to enter the access-list configuration mode).	ip access-list extended <i>name</i>
Add an entry that “points” to the reflexive access list. Add an entry for each reflexive access list <i>name</i> previously defined.	evaluate <i>name</i>

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, and 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 will be evaluated.

If the extended named IP access list you just specified has never been applied to the interface, you must also complete this next task:

Apply the extended named IP access list to the interface, in interface configuration mode:

Task	Command
External interface: apply the extended access list to the interface’s inbound traffic. or Internal interface: apply the extended access list to the interface’s outbound traffic.	ip access-group <i>name</i> in or ip access-group <i>name</i> out

Set a Global Timeout Value (Optional)

Reflexive access list entries expire after no packets in the session have been detected for a certain length of time (the “timeout” period). You can specify the timeout for a particular reflexive access list when you define the reflexive access list. But if you do not specify the timeout for a given reflexive access list, the list will use the global timeout value instead.

The global timeout value is 300 seconds by default. But, you can change the global timeout to a different value at any time.

To change the global timeout value, perform the following task in global configuration mode:

Task	Command
Change the global timeout value for temporary reflexive access list entries.	ip reflexive-list timeout <i>seconds</i>

Reflexive Access Lists Configuration Examples

There are two examples in this section:

- External Interface Configuration Example
- Internal Interface Configuration Example

External Interface Configuration Example

This example has reflexive access lists configured for an external interface, for a topology similar to the one in Figure 7 (shown earlier in this chapter).

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

Define the interface where the session-filtering configuration is to be applied:

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

Apply access lists to the interface, for inbound traffic and for outbound traffic:

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

Define the outbound access list. This is the access list that evaluates all outbound traffic on interface Serial 1.

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

Define the reflexive access list *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 will be automatically created in the reflexive access list *tcptraffic*.

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

Define the inbound access list. This is the access list that evaluates all inbound traffic on interface Serial 1.

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

Define the inbound access list entries. This example shows BGP and Enhanced IGRP running on the interface. Also, no ICMP traffic is permitted. The last entry points to the reflexive access list. If a packet does not match the first three entries, the packet will be evaluated against all the entries in the reflexive access list *tcptraffic*.

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

Define 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 is no TCP traffic that is part of an established session, the corresponding reflexive access list entry will be removed.

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

This is what the example configuration looks like:

```
interface Serial 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 bgp any any
  permit eigrp any any
  deny icmp any any
  evaluate tcptraffic
!
```

With this configuration, before any TCP sessions have been initiated the **show access-list EXEC** command displays the following:

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

Notice that the reflexive access list does not appear in this output. This is because before any TCP sessions have been initiated, no traffic has triggered the reflexive access list, and the list is empty (has no entries). When empty, reflexive access lists do not show up in **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 EXEC** command displays the following:

```
Extended IP access list inboundfilters
  permit bgp any any (2 matches)
  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)
```

Notice that the reflexive access list *tcptraffic* now appears, and displays the temporary entry generated when the Telnet session initiated with an outbound packet.

Internal Interface Configuration Example

This is an example configuration for reflexive access lists configured for an internal interface. This example has a topology similar to the one in Figure 8 (shown earlier in this chapter).

This example is similar to the previous example; the only difference between this example and the previous example is that the entries for the outbound and inbound access lists are swapped. Please refer to the previous example for more details and descriptions.

```
interface Ethernet 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 bgp any any
    permit eigrp any any
    deny icmp any any
    evaluate tcptraffic
  !
  ip access-list extended inboundfilters
    permit tcp any any reflect tcptraffic
  !
```

