



CHAPTER 1

Configuring Multicast Routing

This chapter describes how to configure the ASA to use the multicast routing protocol and includes the following sections:

- [Information About Multicast Routing, page 1-1](#)
- [Licensing Requirements for Multicast Routing, page 1-3](#)
- [Guidelines and Limitations, page 1-3](#)
- [Enabling Multicast Routing, page 1-3](#)
- [Customizing Multicast Routing, page 1-4](#)
- [Configuration Example for Multicast Routing, page 1-15](#)
- [Additional References, page 1-15](#)
- [Feature History for Multicast Routing, page 1-16](#)

Information About Multicast Routing

Multicast routing is a bandwidth-conserving technology that reduces traffic by simultaneously delivering a single stream of information to thousands of corporate recipients and homes. Applications that take advantage of multicast routing include videoconferencing, corporate communications, distance learning, and distribution of software, stock quotes, and news.

Multicast routing protocols deliver source traffic to multiple receivers without adding any additional burden on the source or the receivers while using the least network bandwidth of any competing technology. Multicast packets are replicated in the network by Cisco routers enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols resulting in the most efficient delivery of data to multiple receivers possible.

The ASA supports both stub multicast routing and PIM multicast routing. However, you cannot configure both concurrently on a single ASA.



Note

The UDP and non-UDP transports are both supported for multicast routing. However, the non-UDP transport has no FastPath optimization.

This section includes the following topics:

- [Stub Multicast Routing, page 1-2](#)
- [PIM Multicast Routing, page 1-2](#)

- [Multicast Group Concept, page 1-2](#)
- [Clustering, page 1-2](#)

Stub Multicast Routing

Stub multicast routing provides dynamic host registration and facilitates multicast routing. When configured for stub multicast routing, the ASA acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the ASA forwards IGMP messages to an upstream multicast router, which sets up delivery of the multicast data. When configured for stub multicast routing, the ASA cannot be configured for PIM.

The ASA supports both PIM-SM and bidirectional PIM. PIM-SM is a multicast routing protocol that uses the underlying unicast routing information base or a separate multicast-capable routing information base. It builds unidirectional shared trees rooted at a single Rendezvous Point per multicast group and optionally creates shortest-path trees per multicast source.

PIM Multicast Routing

Bi-directional PIM is a variant of PIM-SM that builds bi-directional shared trees connecting multicast sources and receivers. Bi-directional trees are built using a DF election process operating on each link of the multicast topology. With the assistance of the DF, multicast data is forwarded from sources to the Rendezvous Point, and therefore along the shared tree to receivers, without requiring source-specific state. The DF election takes place during Rendezvous Point discovery and provides a default route to the Rendezvous Point.

**Note**

If the ASA is the PIM RP, use the untranslated outside address of the ASA as the RP address.

Multicast Group Concept

Multicast is based on the concept of a group. An arbitrary group of receivers expresses an interest in receiving a particular data stream. This group does not have any physical or geographical boundaries—the hosts can be located anywhere on the Internet. Hosts that are interested in receiving data flowing to a particular group must join the group using IGMP. Hosts must be a member of the group to receive the data stream.

Multicast Addresses

Multicast addresses specify an arbitrary group of IP hosts that have joined the group and want to receive traffic sent to this group.

Clustering

Multicast routing supports clustering. In Layer 2 clustering, the master unit sends all multicast routing packets and data packets until fast-path forwarding is established. After fast-path forwarding is established, slave units may forward multicast data packets. All data flows are full flows. Stub forwarding flows are also supported. Because only one unit receives multicast packets in Layer 2

clustering, redirection to the master unit is common. In Layer 3 clustering, units do not act independently. All data and routing packets are processed and forwarded by the master unit. Slave units drop all packets that have been sent.

For more information about clustering, see [Chapter 1, “Configuring a Cluster of ASAs.”](#)

Licensing Requirements for Multicast Routing

The following table shows the licensing requirements for this feature:

Model	License Requirement
All models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single context mode. In multiple context mode, unshared interfaces and shared interfaces are not supported.

Firewall Mode Guidelines

Supported only in routed firewall mode. Transparent firewall mode is not supported.

IPv6 Guidelines

Does not support IPv6.

Additional Guidelines

In clustering, for IGMP and PIM, this feature is only supported on the master unit.

Enabling Multicast Routing

Enabling multicast routing lets you enable multicast routing on the ASA. Enabling multicast routing enables IGMP and PIM on all interfaces by default. IGMP is used to learn whether members of a group are present on directly attached subnets. Hosts join multicast groups by sending IGMP report messages. PIM is used to maintain forwarding tables to forward multicast datagrams.



Note

Only the UDP transport layer is supported for multicast routing.

To enable multicast routing, enter the following command:

Command	Purpose
<code>multicast-routing</code>	Enables multicast routing.
Example: <code>hostname (config)# multicast-routing</code>	The number of entries in the multicast routing tables are limited by the amount of RAM on the ASA.

Table 1-1 lists the maximum number of entries for specific multicast tables based on the amount of RAM on the ASA. Once these limits are reached, any new entries are discarded.

Table 1-1 Entry Limits for Multicast Tables

Table	16 MB	128 MB	128+ MB
MFIB	1000	3000	5000
IGMP Groups	1000	3000	5000
PIM Routes	3000	7000	12000

Customizing Multicast Routing

This section describes how to customize multicast routing and includes the following topics:

- [Configuring Stub Multicast Routing and Forwarding IGMP Messages, page 1-4](#)
- [Configuring a Static Multicast Route, page 1-5](#)
- [Configuring IGMP Features, page 1-5](#)
- [Configuring PIM Features, page 1-10](#)
- [Configuring a Bidirectional Neighbor Filter, page 1-13](#)
- [Configuring a Multicast Boundary, page 1-14](#)

Configuring Stub Multicast Routing and Forwarding IGMP Messages



Note

Stub multicast routing and PIM are not supported concurrently.

An ASA acting as the gateway to the stub area does not need to participate in PIM. Instead, you can configure it to act as an IGMP proxy agent and forward IGMP messages from hosts connected on one interface to an upstream multicast router on another interface. To configure the ASA as an IGMP proxy agent, forward the host join and leave messages from the stub area interface to an upstream interface.

To forward the host join and leave messages, enter the following command from the interface attached to the stub area:

Command	Purpose
<pre>igmp forward interface <i>if_name</i></pre> <p>Example: <pre>hostname(config-if)# igmp forward interface <i>interface1</i></pre></p>	Configures stub multicast routing and forwards IGMP messages.

Configuring a Static Multicast Route

Configuring static multicast routes lets you separate multicast traffic from unicast traffic. For example, when a path between a source and destination does not support multicast routing, the solution is to configure two multicast devices with a GRE tunnel between them and to send the multicast packets over the tunnel.

When using PIM, the ASA expects to receive packets on the same interface where it sends unicast packets back to the source. In some cases, such as bypassing a route that does not support multicast routing, you may want unicast packets to take one path and multicast packets to take another.

Static multicast routes are not advertised or redistributed.

To configure a static multicast route or a static multicast route for a stub area, enter one of the following commands:

Command	Purpose
<pre>mroute <i>src_ip src_mask</i> {<i>input_if_name</i> <i>rpf_neighbor</i>} [<i>distance</i>]</pre> <p>Example: <pre>hostname(config)# mroute <i>src_ip src_mask</i> {<i>input_if_name</i> <i>rpf_neighbor</i>} [<i>distance</i>]</pre></p>	Configures a static multicast route.
<pre>mroute <i>src_ip src_mask</i> <i>input_if_name</i> [<i>dense output_if_name</i>] [<i>distance</i>]</pre> <p>Example: <pre>hostname(config)# mroute <i>src_ip src_mask</i> <i>input_if_name</i> [<i>dense output_if_name</i>] [<i>distance</i>]</pre></p>	Configures a static multicast route for a stub area. The dense <i>output_if_name</i> keyword and argument pair is only supported for stub multicast routing.

Configuring IGMP Features

IP hosts use the Internet Group Management Protocol (IGMP) to report their group memberships to directly connected multicast routers.

IGMP is used to dynamically register individual hosts in a multicast group on a particular LAN. Hosts identify group memberships by sending IGMP messages to their local multicast router. Under IGMP, routers listen to IGMP messages and periodically send out queries to discover which groups are active or inactive on a particular subnet.

IGMP uses group addresses (Class D IP address) as group identifiers. Host group address can be in the range of 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is never assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.

When you enable multicast routing on the ASA, IGMP Version 2 is automatically enabled on all interfaces.

**Note**

Only the **no igmp** command appears in the interface configuration when you use the **show run** command. If the **multicast-routing** command appears in the device configuration, then IGMP is automatically enabled on all interfaces.

This section describes how to configure optional IGMP setting on a per-interface basis and includes the following topics:

- [Disabling IGMP on an Interface, page 1-6](#)
- [Configuring IGMP Group Membership, page 1-7](#)
- [Configuring a Statically Joined IGMP Group, page 1-7](#)
- [Controlling Access to Multicast Groups, page 1-8](#)
- [Limiting the Number of IGMP States on an Interface, page 1-8](#)
- [Modifying the Query Messages to Multicast Groups, page 1-8](#)
- [Changing the IGMP Version, page 1-9](#)

Disabling IGMP on an Interface

You can disable IGMP on specific interfaces. This information is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the ASA from sending host query messages on that interface.

To disable IGMP on an interface, enter the following command:

Command	Purpose
<code>no igmp</code>	Disables IGMP on an interface. To reenable IGMP on an interface, use the igmp command.
Example: <code>hostname(config-if)# no igmp</code>	

**Note**

Only the **no igmp** command appears in the interface configuration.

Configuring IGMP Group Membership

You can configure the ASA to be a member of a multicast group. Configuring the ASA to join a multicast group causes upstream routers to maintain multicast routing table information for that group and keep the paths for that group active.



Note

If you want to forward multicast packets for a specific group to an interface without the ASA accepting those packets as part of the group, see the [“Configuring a Statically Joined IGMP Group” section on page 1-7](#).

To have the ASA join a multicast group, enter the following command:

Command	Purpose
<pre>igmp join-group group-address</pre> <p>Example: <pre>hostname(config-if)# igmp join-group mcast-group</pre></p>	<p>Configures the ASA to be a member of a multicast group.</p> <p>The <i>group-address</i> argument is the IP address of the group.</p>

Configuring a Statically Joined IGMP Group

Sometimes a group member cannot report its membership in the group because of some configuration, or there may be no members of a group on the network segment. However, you still want multicast traffic for that group to be sent to that network segment. You can have multicast traffic for that group sent to the segment by configuring a statically joined IGMP group.

Enter the **igmp static-group** command. The ASA does not accept the multicast packets, but instead forwards them to the specified interface.

To configure a statically joined multicast group on an interface, enter the following command:

Command	Purpose
<pre>igmp static-group</pre> <p>Example: <pre>hostname(config-if)# igmp static-group group-address</pre></p>	<p>Configures the ASA statically to join a multicast group on an interface.</p> <p>The <i>group-address</i> argument is the IP address of the group.</p>

Controlling Access to Multicast Groups

To control the multicast groups that hosts on the ASA interface can join, perform the following steps:

Detailed Steps

Command	Purpose
Step 1 Do one of the following to create a standard or extended access list:	
<pre>access-list name standard [permit deny] ip_addr mask</pre> <p>Example: hostname(config)# access-list acl1 standard permit 192.52.662.25 </p>	Creates a standard access list for the multicast traffic. You can create more than one entry for a single access list. You can use extended or standard access lists. The <i>ip_addr mask</i> argument is the IP address of the multicast group being permitted or denied.
<pre>access-list name extended [permit deny] protocol src_ip_addr src_mask dst_ip_addr dst_mask</pre> <p>Example: hostname(config)# access-list acl2 extended permit protocol src_ip_addr src_mask dst_ip_addr dst_mask </p>	Creates an extended access list. The <i>dst_ip_addr</i> argument is the IP address of the multicast group being permitted or denied.
Step 2 <pre>igmp access-group acl</pre> <p>Example: hostname(config-if)# igmp access-group acl </p>	Applies the access list to an interface. The <i>acl</i> argument is the name of a standard or extended IP access list.

Limiting the Number of IGMP States on an Interface

You can limit the number of IGMP states resulting from IGMP membership reports on a per-interface basis. Membership reports exceeding the configured limits are not entered in the IGMP cache, and traffic for the excess membership reports is not forwarded.

To limit the number of IGMP states on an interface, enter the following command:

Command	Purpose
<pre>igmp limit number</pre> <p>Example: hostname(config-if)# igmp limit 50 </p>	Limits the number of IGMP states on an interface. Valid values range from 0 to 500, with 500 being the default value. Setting this value to 0 prevents learned groups from being added, but manually defined memberships (using the igmp join-group and igmp static-group commands) are still permitted. The no form of this command restores the default value.

Modifying the Query Messages to Multicast Groups



Note The **igmp query-timeout** and **igmp query-interval** commands require IGMP Version 2.

The ASA sends query messages to discover which multicast groups have members on the networks attached to the interfaces. Members respond with IGMP report messages indicating that they want to receive multicast packets for specific groups. Query messages are addressed to the all-systems multicast group, which has an address of 224.0.0.1, with a time-to-live value of 1.

These messages are sent periodically to refresh the membership information stored on the ASA. If the ASA discovers that there are no local members of a multicast group still attached to an interface, it stops forwarding multicast packet for that group to the attached network, and it sends a prune message back to the source of the packets.

By default, the PIM designated router on the subnet is responsible for sending the query messages. By default, they are sent once every 125 seconds.

When changing the query response time, by default, the maximum query response time advertised in IGMP queries is 10 seconds. If the ASA does not receive a response to a host query within this amount of time, it deletes the group.

To change the query interval, query response time, and query timeout value, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<code>igmp query-interval <i>seconds</i></code> Example: hostname(config-if)# igmp query-interval 30	Sets the query interval time in seconds. Valid values range from 0 to 500; 125 is the default value. If the ASA does not hear a query message on an interface for the specified timeout value (by default, 255 seconds), then the ASA becomes the designated router and starts sending the query messages.
Step 2	<code>igmp query-timeout <i>seconds</i></code> Example: hostname(config-if)# igmp query-timeout 30	Changes the timeout value of the query. Valid values range from 0 to 500; 225 is the default value.
Step 3	<code>igmp query-max-response-time <i>seconds</i></code> Example: hostname(config-if)# igmp query-max-response-time 30	Changes the maximum query response time.

Changing the IGMP Version

By default, the ASA runs IGMP Version 2, which enables several additional features such as the **igmp query-timeout** and **igmp query-interval** commands.

All multicast routers on a subnet must support the same version of IGMP. The ASA does not automatically detect Version 1 routers and switch to Version 1. However, a mix of IGMP Version 1 and 2 hosts on the subnet works; the ASA running IGMP Version 2 works correctly when IGMP Version 1 hosts are present.

To control which version of IGMP is running on an interface, enter the following command:

Command	Purpose
<code>igmp version {1 2}</code>	Controls the version of IGMP that you want to run on the interface.
Example: <code>hostname(config-if)# igmp version 2</code>	

Configuring PIM Features

Routers use PIM to maintain forwarding tables for forwarding multicast diagrams. When you enable multicast routing on the ASA, PIM and IGMP are automatically enabled on all interfaces.



Note

PIM is not supported with PAT. The PIM protocol does not use ports, and PAT only works with protocols that use ports.

This section describes how to configure optional PIM settings and includes the following topics:

- [Enabling and Disabling PIM on an Interface, page 1-10](#)
- [Configuring a Static Rendezvous Point Address, page 1-11](#)
- [Configuring the Designated Router Priority, page 1-11](#)
- [Configuring and Filtering PIM Register Messages, page 1-12](#)
- [Configuring PIM Message Intervals, page 1-12](#)
- [Filtering PIM Neighbors, page 1-12](#)

Enabling and Disabling PIM on an Interface

You can enable or disable PIM on specific interfaces. To enable or disable PIM on an interface, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<code>pim</code> Example: <code>hostname(config-if)# pim</code>	Enables or reenables PIM on a specific interface.
Step 2	<code>no pim</code> Example: <code>hostname(config-if)# no pim</code>	Disables PIM on a specific interface.



Note

Only the **no pim** command appears in the interface configuration.

Configuring a Static Rendezvous Point Address

All routers within a common PIM sparse mode or bidir domain require knowledge of the PIM RP address. The address is statically configured using the **pim rp-address** command.



Note

The ASA does not support Auto-RP or PIM BSR. You must use the **pim rp-address** command to specify the RP address.

You can configure the ASA to serve as RP to more than one group. The group range specified in the access list determines the PIM RP group mapping. If an access list is not specified, then the RP for the group is applied to the entire multicast group range (224.0.0.0/4).

To configure the address of the PIM RP, enter the following command:

Command	Purpose
<p>pim rp-address <i>ip_address</i> [<i>acl</i>] [<i>bidir</i>]</p> <p>Example: <pre>hostname(config)# pim rp-address 10.86.75.23 [acl1] [bidir]</pre></p>	<p>Enables or reenables PIM on a specific interface.</p> <p>The <i>ip_address</i> argument is the unicast IP address of the router assigned to be a PIM RP.</p> <p>The <i>acl</i> argument is the name or number of a standard access list that defines with which multicast groups the RP should be used. Do not use a host ACL with this command.</p> <p>Excluding the bidir keyword causes the groups to operate in PIM sparse mode.</p>



Note

The ASA always advertises the bidirectional capability in the PIM hello messages, regardless of the actual bidirectional configuration.

Configuring the Designated Router Priority

The DR is responsible for sending PIM register, join, and prune messages to the RP. When there is more than one multicast router on a network segment, selecting the DR is based on the DR priority. If multiple devices have the same DR priority, then the device with the highest IP address becomes the DR.

By default, the ASA has a DR priority of 1. To change this value, enter the following command:

Command	Purpose
<p>pim dr-priority <i>num</i></p> <p>Example: <pre>hostname(config-if)# pim dr-priority 500</pre></p>	<p>Changes the designated router priority.</p> <p>The <i>num</i> argument can be any number ranging from 1 to 4294967294.</p>

Configuring and Filtering PIM Register Messages

When the ASA is acting as an RP, you can restrict specific multicast sources from registering with it to prevent unauthorized sources from registering with the RP. The Request Filter pane lets you define the multicast sources from which the ASA will accept PIM register messages.

To filter PIM register messages, enter the following command:

Command	Purpose
<pre>pim accept-register {list <i>acl</i> route-map map-name}</pre> <p>Example: <pre>hostname(config)# pim accept-register {list <i>acl1</i> route-map <i>map2</i>}</pre></p>	<p>Configures the ASA to filter PIM register messages.</p> <p>In the example, the ASA filters PIM register messages <i>acl1</i> and route map <i>map2</i>.</p>

Configuring PIM Message Intervals

Router query messages are used to select the PIM DR. The PIM DR is responsible for sending router query messages. By default, router query messages are sent every 30 seconds. Additionally, every 60 seconds, the ASA sends PIM join or prune messages.

To change these intervals, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>pim hello-interval <i>seconds</i></pre> <p>Example: <pre>hostname(config-if)# pim hello-interval 60</pre></p>	<p>Sends router query messages.</p> <p>Valid values for the <i>seconds</i> argument range from 1 to 3600 seconds.</p>
Step 2	<pre>pim join-prune-interval <i>seconds</i></pre> <p>Example: <pre>hostname(config-if)# pim join-prune-interval 60</pre></p>	<p>Changes the amount of time (in seconds) that the ASA sends PIM join or prune messages.</p> <p>Valid values for the <i>seconds</i> argument range from 10 to 600 seconds.</p>

Filtering PIM Neighbors

You can define the routers that can become PIM neighbors. By filtering the routers that can become PIM neighbors, you can do the following:

- Prevent unauthorized routers from becoming PIM neighbors.
- Prevent attached stub routers from participating in PIM.

To define neighbors that can become a PIM neighbor, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>access-list pim_nbr deny router-IP_addr PIM neighbor</pre> <p>Example: <pre>hostname(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255</pre></p>	<p>Uses a standard access list to define the routers that you want to have participate in PIM.</p> <p>In the example, the following access list, when used with the pim neighbor-filter command, prevents the 10.1.1.1 router from becoming a PIM neighbor.</p>
Step 2	<pre>pim neighbor-filter pim_nbr</pre> <p>Example: <pre>hostname(config)# interface GigabitEthernet0/3 hostname(config-if)# pim neighbor-filter pim_nbr</pre></p>	<p>Filters neighbor routers.</p> <p>In the example, the 10.1.1.1 router is prevented from becoming a PIM neighbor on interface GigabitEthernet0/3.</p>

Configuring a Bidirectional Neighbor Filter

The Bidirectional Neighbor Filter pane shows the PIM bidirectional neighbor filters, if any, that are configured on the ASA. A PIM bidirectional neighbor filter is an ACL that defines the neighbor devices that can participate in the DF election. If a PIM bidirectional neighbor filter is not configured for an interface, then there are no restrictions. If a PIM bidirectional neighbor filter is configured, only those neighbors permitted by the ACL can participate in the DF election process.

When a PIM bidirectional neighbor filter configuration is applied to the ASA, an ACL appears in the running configuration with the name *interface-name_multicast*, in which the *interface-name* is the name of the interface to which the multicast boundary filter is applied. If an ACL with that name already exists, a number is appended to the name (for example, *inside_multicast_1*). This ACL defines which devices can become PIM neighbors of the ASA.

Bidirectional PIM allows multicast routers to keep reduced state information. All of the multicast routers in a segment must be bidirectionally enabled for *bidir* to elect a DF.

The PIM bidirectional neighbor filters enable the transition from a sparse-mode-only network to a *bidir* network by letting you specify the routers that should participate in the DF election, while still allowing all routers to participate in the sparse-mode domain. The *bidir*-enabled routers can elect a DF from among themselves, even when there are non-*bidir* routers on the segment. Multicast boundaries on the non-*bidir* routers prevent PIM messages and data from the *bidir* groups from leaking in or out of the *bidir* subset cloud.

When a PIM bidirectional neighbor filter is enabled, the routers that are permitted by the ACL are considered to be bidirectionally capable. Therefore, the following is true:

- If a permitted neighbor does not support *bidir*, then the DF election does not occur.
- If a denied neighbor supports *bidir*, then the DF election does not occur.
- If a denied neighbor does not support *bidir*, the DF election can occur.

To define the neighbors that can become a PIM bidirectional neighbor filter, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	<pre>access-list pim_nbr deny router-IP_addr PIM neighbor</pre> <p>Example: hostname(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255 </p>	<p>Uses a standard access list to define the routers that you want to have participate in PIM.</p> <p>In the example, the following access list, when used with the pim neighbor-filter command, prevents the 10.1.1.1 router from becoming a PIM neighbor.</p>
Step 2	<pre>pim bidirectional-neighbor-filter pim_nbr</pre> <p>Example: hostname(config)# interface GigabitEthernet0/3 hostname(config-if)# pim bidirectional neighbor-filter pim_nbr </p>	<p>Filters neighbor routers.</p> <p>In the example, the 10.1.1.1 router is prevented from becoming a PIM bidirectional neighbor on interface GigabitEthernet0/3.</p>

Configuring a Multicast Boundary

Address scoping defines domain boundaries so that domains with RPs that have the same IP address do not leak into each other. Scoping is performed on the subnet boundaries within large domains and on the boundaries between the domain and the Internet.

You can set up an administratively scoped boundary on an interface for multicast group addresses by entering the **multicast boundary** command. IANA has designated the multicast address range from 239.0.0.0 to 239.255.255.255 as the administratively scoped addresses. This range of addresses can be reused in domains administered by different organizations. The addresses would be considered local, not globally unique.

A standard ACL defines the range of affected addresses. When a boundary is set up, no multicast data packets are allowed to flow across the boundary from either direction. The boundary allows the same multicast group address to be reused in different administrative domains.

You can configure, examine, and filter Auto-RP discovery and announcement messages at the administratively scoped boundary by entering the **filter-autorp** keyword. Any Auto-RP group range announcements from the Auto-RP packets that are denied by the boundary ACL are removed. An Auto-RP group range announcement is permitted and passed by the boundary only if all addresses in the Auto-RP group range are permitted by the boundary ACL. If any address is not permitted, the entire group range is filtered and removed from the Auto-RP message before the Auto-RP message is forwarded.

To configure a multicast boundary, enter the following command:

Command	Purpose
<pre>multicast boundary acl [filter-autorp]</pre> <p>Example: hostname(config-if)# multicast boundary acl1 [filter-autorp] </p>	<p>Configures a multicast boundary.</p>

Configuration Example for Multicast Routing

The following example shows how to enable and configure multicast routing with various optional processes:

Step 1 Enable multicast routing:

```
hostname (config) # multicast-routing
```

Step 2 Configure a static multicast route:

```
hostname (config) # mroute src_ip src_mask {input_if_name | rpf_neighbor} [distance]  
hostname (config) # exit
```

Step 3 Configure the ASA to be a member of a multicast group:

```
hostname (config) # interface  
hostname (config-if) # igmp join-group group-address
```

Additional References

For additional information related to routing, see the following sections:

- [Related Documents, page 1-16](#)
- [RFCs, page 1-16](#)

Related Documents

Related Topic	Document Title
Technical details about the IGMP and multicast routing standards used for implementing the SMR feature	IETF draft-ietf-idmr-igmp-proxy-01.txt

RFCs

RFC	Title
RFC 2113	IP Router Alert Option
RFC 2236	IGMPv2
RFC 2362	PIM-SM
RFC 2588	IP Multicast and Firewalls

Feature History for Multicast Routing

[Table 1-2](#) lists each feature change and the platform release in which it was implemented.

Table 1-2 Feature History for Multicast Routing

Feature Name	Platform Releases	Feature Information
Multicast routing support	7.0(1)	Support was added for multicast routing data, authentication, and redistribution and monitoring of routing information using the multicast routing protocol. We introduced the multicast-routing command.
Clustering support	9.0(1)	Support was added for clustering. We introduced the following commands: debug mfib cluster , show mfib cluster .