CHAPTER 24

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

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 delivers 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.

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 bi-directional 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.

Licensing Requirements for Multicast Routing

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, shared interfaces are not supported.

Firewall Mode Guidelines

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

IPv6 Guidelines

Does not support IPv6.

Enabling Multicast Routing

Enabling multicast routing lets the ASA forward multicast packets. Enabling multicast routing automatically enables PIM and IGMP on all interfaces.

To enable multicast routing, perform the following step:

Detailed Steps

Command	Purpose
multicast-routing **Example:** hostname(config)# multicast-routing	This step enables multicast routing. The number of entries in the multicast routing tables are limited by the amount of RAM on the system.

Table 24-1 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 24-1 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

Note Stub Multicast Routing and PIM are not supported concurrently.

A 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. 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, perform the following step from the interface attached to the stub area:

Detailed Steps

Command	Purpose
igmp forward interface *if_name* **Example:** `hostname(config-if)# igmp forward interface interface1`	This step configures stub multicast routing.

Configuring a Static Multicast Route

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, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	Do one of the following to configure a static multicast route or a static multicast route for a stub area.	
	mroute *src_ip src_mask* {*input_if_name* | *rpf_neighbor*} [*distance*] **Example:** hostname(config)# mroute *src_ip src_mask* {*input_if_name* | *rpf_neighbor*} [*distance*]	This step configures a static multicast route.
	mroute *src_ip src_mask input_if_name* [**dense** *output_if_name*] [*distance*] **Example:** hostname(config)# mroute *src_ip src_mask input_if_name* [dense *output_if_name*] [*distance*]	This step configures a static multicast route for a stub area. The **dense** *output_if_name* keyword and argument pair is only supported for stub multicast routing.

Configuring IGMP Features

IP hosts use Internet Group Management Protocol, or 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 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. This section includes the following topics:

- Disabling IGMP on an Interface, page 24-22
- Configuring IGMP Group Membership, page 24-22
- Configuring a Statically Joined IGMP Group, page 24-22
- Controlling Access to Multicast Groups, page 24-23
- Limiting the Number of IGMP States on an Interface, page 24-23

- Modifying the Query Messages to Multicast Groups, page 24-24
- Changing the IGMP Version, page 24-25

Disabling IGMP on an Interface

You can disable IGMP on specific interfaces. This is useful if you know that you do not have any 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, perform the following steps:

Detailed Steps

Command	Purpose
no igmp **Example:** `hostname(config-if)# no igmp`	This step disables IGMP on an interface. To reenable IGMP on an interface, do the following: `hostname(config-if)# igmp`

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.

To have the ASA join a multicast group, perform the following steps:

Detailed Steps

Command	Purpose
igmp join-group *group-address* **Example:** `hostname(config-if)# igmp join-group` *mcast-group*	This step configures the ASA to be a member of a multicast group. The *group-address* is the IP address of the group.

Configuring a Statically Joined IGMP Group

Sometimes a group member cannot report its membership in the group, or there may be no members of a group on the network segment, but 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 in one of two ways:

- Using the **igmp join-group** command (see Configuring IGMP Group Membership, page 24-22). This causes the ASA to accept and to forward the multicast packets.

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

To configure a statically joined multicast group on an interface,perform the following steps:

Detailed Steps

Command	Purpose

`igmp static-group`
`Example:`
`hostname(config-if)#` **`igmp static-group`** *`group-address`*
	This step configures the ASA statistically join a multicast group on an interface.
	The *`group-address`* is the IP address of the group.

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 to create a standard or extended access list.	

	`access-list` *`name`* **`standard`** `[`**`permit`** `|` **`deny`**`]` *`ip_addr mask`*
	`Example:`
	`hostname(config)# access-list` *`acl1`* `standard permit` *`192.52.662.25`*
		This step 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 *ip_addr mask* argument is the IP address of the multicast group being permitted or denied.

	`access-list` *`name`* **`extended`** `[`**`permit`** `|` **`deny`**`]` *`protocol`* `src_ip_addr src_mask` *`dst_ip_addr dst_mask`*
	`Example:`
	`hostname(config)# access-list` *`acl2`* `extended permit` *`protocol`* `src_`*`ip_addr`* `src_`*`mask dst_ip_addr dst_mask`*
		This step creates an extended access list.
		The *dst_ip_addr* argument is the IP address of the multicast group being permitted or denied.

Step 2	**`igmp access-group`** *`acl`*
	`Example:`
	`hostname(config-if)# igmp access-group` *`acl`*
		Apply the access list to an interface.
		The *acl* 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, perform the following steps:

Detailed Steps

Command	Purpose
igmp limit *number* **Example:** `hostname(config-if)# igmp limit 50`	This limit 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	**igmp query-interval** *seconds* **Example:** `hostname(config-if)# igmp query-interval 30`	To set the query interval time in seconds. Valid values range from 0 to 500, with 125 being 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.

	Command	Purpose
Step 2	**igmp query-timeout** *seconds* **Example:** hostname(config-if)# **igmp query-timeout** *30*	To change this timeout value of the query. Valid values range from 0 to 500, with 225 being the default value.
Step 3	**igmp query-max-response-time** *seconds* **Example:** hostname(config-if)# igmp query-max-response-time *30*	To change 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, perform the following steps:

Detailed Steps

Command	Purpose
igmp version {1 | 2} **Example:** hostname(config-if)# igmp version 2	This step controls which version of IGMP you want to run on the interface.

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. This section includes the following topics:

- Enabling and Disabling PIM on an Interface, page 24-26
- Configuring a Static Rendezvous Point Address, page 24-26
- Configuring the Designated Router Priority, page 24-27
- Filtering PIM Register Messages, page 24-28
- Configuring PIM Message Intervals, page 24-28
- Configuring a Multicast Boundary, page 24-28

- Filtering PIM Neighbors, page 24-29
- Supporting Mixed Bidirectional/Sparse-Mode PIM Networks, page 24-29

Enabling and Disabling PIM on an Interface

You can disable PIM on specific interfaces. To disable PIM on an interface, use the following steps

Detailed Steps

	Command	Purpose

Step 1	**`pim`**
	Example:
	`hostname(config-if)# pim`
		This step enables or reenables PIM on a specific interface.

Step 2	**`no pim`**
	Example:
	`hostname(config-if)# no pim`
		This step 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 PR, use the following step:

Detailed Steps

Command	Purpose

pim rp-address *ip_address* [*acl*] [**bidir**]
Example:
`hostname(config)# pim rp-address`
`ip_address` `[`*`acl`*`] [bidir]`
	This step enables or reenables PIM on a specific interface.
	The *ip_address* argument is the unicast IP address of the router to be a PIM RP.
	The *acl* argument is the name or number of a standard access list that defines which multicast groups the RP should be used with. Do not use a host ACL with this command.
	Excluding the **bidir** keyword causes the groups to operate in PIM sparse mode.

Note The ASA always advertises the bidir capability in the PIM hello messages regardless of the actual bidir configuration.

Configuring the Designated Router Priority

The DR is responsible for sending PIM register, join, and prune messaged to the RP. When there is more than one multicast router on a network segment, there is an election process to select the DR based on 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. You can change this value by performing this step:

Detailed Steps

Command	Purpose
pim dr-priority *num* **Example:** `hostname(config-if)# pim dr-priority` *`500`*	This step changes the designated router priority. The *num* argument can be any number from 1 to 4294967294.

Filtering PIM Register Messages

You can configure the ASA to filter PIM register messages. To filter PIM register messages, perform the following step:

Detailed Steps

Command	Purpose
pim accept-register {**list** *acl* | **route-map** *map-name*} **Example:** `hostname(config)# pim accept-register {list acl	route-map map-name}`	This step configure the ASA to filter PIM register messages.

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/prune messages. To change these intervals, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**pim hello-interval** *seconds* **Example:** `hostname(config-if)# pim hello-interval 60`	This step sends router query messages. Valid values for the *seconds* argument range from 1 to 3600 seconds.
Step 2	**pim join-prune-interval** *seconds* **Example:** `hostname(config-if)# pim join-prune-interval 60`	This step changes the amount of time (in seconds) that the ASA sends PIM join/prune messages. Valid values for the *seconds* argument range from 10 to 600 seconds

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 using the **multicast boundary** command. IANA has designated the multicast address range 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. They would be considered local, not globally unique.

A standard ACL defines the range of addresses affected. 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 the **filter-autorp** keyword to examine and filter Auto-RP discovery and announcement messages at the administratively scoped boundary. Any Auto-RP group range announcements from the Auto-RP packets that are denied by the boundary access control list (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, perform the following step:

Detailed Steps

Command	Purpose

`multicast boundary acl [filter-autorp]`
Example:
`hostname(config-if)# multicast boundary acl [filter-autorp]`
	This step configures a multicast boundary.

Filtering PIM Neighbors

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

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

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

Detailed Steps

	Command	Purpose

Step 1	`access-list pim_nbr deny router-IP_addr PIM neighbor`
	Example:
	`hostname(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255`
		This step uses the **access-list** command to define a standard access list defines the routers you want to participate in PIM.
		In this 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:

Step 2	`pim neighbor-filter pim_nbr`
	Example:
	`hostname(config)# interface GigabitEthernet0/3`
	`hostname(config-if)# pim neighbor-filter pim_nbr`
		Use the **pim neighbor-filter** command on an interface to filter the neighbor routers.
		In this example, the 10.1.1.1 router is prevented from becoming a PIM neighbor on interface GigabitEthernet0/3.

Supporting Mixed Bidirectional/Sparse-Mode PIM Networks

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

Bidirectional PIM enables the transition from a sparse-mode-only network to a bidir network by letting you specify the routers that should participate in 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 bidirectional PIM is enabled, the routers that are permitted by the ACL are considered to be bidir-capable. Therefore:

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

To control which neighbors can participate in the DF election, perform the following steps:

Detailed Steps

	Command	Purpose
Step 1	**access-list pim_bidir deny any** **Example:** hostname(config)# access-list pim_bidir permit 10.1.1.1 255.255.255.255 hostname(config)# access-list pim_bidir permit 10.1.1.2 255.255.255.255 hostname(config)# access-list pim_bidir deny any	This step uses the **access-list** command to define a standard access list defines the routers you want to participate in in the DF election and denies all others. In this example, the following access list permits the routers at 10.1.1.1 and 10.2.2.2 to participate in the DF election and denies all others.
Step 2	**pim bidir-neighbor-filter pim_bidir** **Example:** hostname(config)# interface GigabitEthernet0/3 hostname(config-if)# pim bidir-neighbor-filter pim_bidir	Enable bidirectional PIM on an interface. This example applies the access list created previous step to the interface GigabitEthernet0/3.

Configuration Example for Multicast Routing

The following example shows how to enable and configure muticastrouting 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 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:

- Related Documents, page 24-31
- RFCs, page 24-31

Related Documents

Related Topic	Document Title
Routing Overview	Information About Routing
How to configure OSPF	Configuring OSPF
How to configure EIGRP	Configuring EIGRP
How to configure RIP	Configuring RIP
How to configure a static or default route	Configuring Static and Default Routes
How to configure a route map	Defining Route Maps

RFCs

The following is list of RFCs from the IETF provide technical details about the IGMP and multicast routing standards used for implementing the SMR feature:

- RFC 2236 IGMPv2
- RFC 2362 PIM-SM
- RFC 2588 IP Multicast and Firewalls
- RFC 2113 IP Router Alert Option
- IETF draft-ietf-idmr-igmp-proxy-01.txt