Multicast Routing

This chapter describes how to configure the Cisco ASA to use the multicast routing protocol.

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 ASA enabled with Protocol Independent Multicast (PIM) and other supporting multicast protocols, which results 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 sparse or bidirectional mode. You must enable PIM on the interfaces participating in IGMP stub multicast routing.

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 (RP) per multicast group and optionally creates shortest-path trees per multicast source.

PIM Multicast Routing

Bidirectional PIM is a variant of PIM-SM that builds bidirectional shared trees connecting multicast sources and receivers. Bidirectional trees are built using a Designated Forwarder (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 (RP), and therefore along the shared tree to receivers, without requiring source-specific state. The DF election takes place during RP discovery and provides a default route to the RP.

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 Spanned EtherChannel clustering, the control unit sends all multicast routing packets and data packets until fast-path forwarding is established. After fast-path forwarding is established, data 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 Spanned EtherChannel clustering, redirection to the control unit is common. In Individual Interface clustering, units do not act independently. All data and routing packets are processed and forwarded by the control unit. Data units drop all packets that have been sent.

For more information about clustering, see ASA Cluster.

Guidelines for Multicast Routing

Context Mode

Supported in single context mode.

Firewall Mode

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

IPv6

Does not support IPv6.

Multicast Group

The range of addresses between 224.0.0.0 and 224.0.0.255 is reserved for the use of routing protocols and other topology discovery or maintenance protocols, such as gateway discovery and group membership reporting. Hence, Internet multicast routing from address range 224.0.0/24 is not supported; IGMP group is not created when enabling multicast routing for the reserved addressess.

Clustering

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

Additional Guidelines

- You must configure an access control rule on the inbound interface to allow traffic to the multicast host, such as 224.1.2.3. However, you cannot specify a destination interface for the rule, or it cannot be applied to multicast connections during initial connection validation.
- PIM/IGMP Multicast routing is not supported on interfaces in a traffic zone.

Enable Multicast Routing

Enabling multicast routing on the ASA, enables IGMP and PIM on all data interfaces by default, but not on the management interface for most models (see Management Slot/Port Interface for interfaces that do not allow through traffic). 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.

To enable multicast routing on the management interface, you must explicitly set a multicast boundary on the management interface.

Note Only the UDP transport layer is supported for multicast routing.

The following list shows the maximum number of entries for specific multicast tables. Once these limits are reached, any new entries are discarded.

- MFIB—30,000
- IGMP Groups—30,000
- PIM Routes—72,000

Procedure

Enable multicast routing:

multicast-routing

Example:

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

The number of entries in the multicast routing tables are limited by the amount of RAM on the ASA.

Customize Multicast Routing

This section describes how to customize multicast routing.

Configure Stub Multicast Routing and Forward IGMP Messages

Note Stub multicast routing is not supported concurrently with PIM sparse and bidirectional modes.

An ASA acting as the gateway to the stub area does not need to participate in PIM sparse mode or bidirectional mode. 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. You must also enable PIM on the interfaces participating in stub mode multicast routing.

Procedure

Configure stub multicast routing and forward IGMP messages:

igmp forward interface if_name

Example:

```
ciscoasa(config-if)# igmp forward interface interface1
```

Configure 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.

Procedure

Step 1 Configure a static multicast route:

mroute src_ip src_mask {input_if_name | *rpf_neighbor*} [distance]

Example:

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

Step 2 Configure a static multicast route for a stub area:

mroute src_ip src_mask input_if_name [**dense** output_if_name] [distance]

Example:

```
ciscoasa(config)# mroute src_ip src_mask input_if_name [dense output_if_name] [distance]
```

The **dense** *output_if_name* keyword and argument pair is only supported for stub multicast routing.

Configure IGMP Features

IP hosts use 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.

This section describes how to configure optional IGMP setting on a per-interface basis.

Disable 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.

Procedure

Disable IGMP on an interface:

no igmp

Example:

```
ciscoasa(config-if)# no igmp
```

To reenable IGMP on an interface, use the **igmp** command.

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

Configure 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 Configure a Statically Joined IGMP Group, on page 6.

Procedure

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

igmp join-group group-address

Example:

```
ciscoasa(config-if)# igmp join-group mcast-group
```

The *group-address* argument is the IP address of the group.

Configure 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.

Procedure

Configure the ASA statically to join a multicast group on an interface:

igmp static-group

Example:

```
ciscoasa(config-if)# igmp static-group group-address
```

The *group-address* argument is the IP address of the group.

Control Access to Multicast Groups

You can control access to multicast groups by using access control lists.

Procedure

Step 1 Create a standard ACL for the multicast traffic:

access-list name **standard** [**permit | deny]** ip_addr mask

Example:

```
ciscoasa(config)# access-list acl1 standard permit 192.52.662.25
```

You can create more than one entry for a single ACL. You can use extended or standard ACLs.

The *ip_addr mask* argument is the IP address of the multicast group being permitted or denied.

Step 2 Create an extended ACL:'

access-list name **extended** [**permit | deny]** protocol src_ip_addr src_mask dst_ip_addr dst_mask

Example:

```
ciscoasa(config)# access-list acl2 extended permit protocol
src_ip_addr src_mask dst_ip_addr dst_mask
```

The *dst_ip_addr* argument is the IP address of the multicast group being permitted or denied.

Step 3 Apply the ACL to an interface:

igmp access-group acl

Example:

```
ciscoasa(config-if)# igmp access-group acl
```

The *acl* argument is the name of a standard or extended IP ACL.

Limit 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.

Procedure

Limit the number of IGMP states on an interface:

igmp limit number

Example:

```
ciscoasa(config-if)# igmp limit 50
```

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.

Modify the Query Messages to Multicast Groups

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.

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

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

Procedure

Step 1 Set the query interval time in seconds:

igmp query-interval *seconds*

Example:

```
ciscoasa(config-if)# igmp query-interval 30
```

Valid values range from 1 to 3600; 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 Change the timeout value of the query:

igmp query-timeout *seconds*

Example:

```
ciscoasa(config-if)# igmp query-timeout 30
```

Valid values range from 60 to 300; 225 is the default value.

Step 3 Change the maximum query response time:

igmp query-max-response-time *seconds*

Valid values range from 1 to 25; 10 is the default value.

Example:

```
ciscoasa(config-if)# igmp query-max-response-time 20
```

Change 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.

Procedure

Control the version of IGMP that you want to run on the interface:

igmp version {**1** | **2**}

Example:

```
ciscoasa(config-if)# igmp version 2
```

Configure PIM Features

Routers use PIM to maintain forwarding tables to use 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.

Enable and Disable PIM on an Interface

You can enable or disable PIM on specific interfaces.

Procedure

Step 1 Enable or reenable PIM on a specific interface:

pim

Example:

```
ciscoasa(config-if)# pim
```

Step 2 Disable PIM on a specific interface:

no pim

Example:

```
ciscoasa(config-if)# no pim
```

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

Configure 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. 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 ACL determines the PIM RP group mapping. If an ACL is not specified, then the RP for the group is applied to the entire multicast group range (224.0.0.0/4).

Procedure

Enable or reenable PIM on a specific interface:

pim rp-address ip_address [acl] [**bidir**]

The *ip_address* argument is the unicast IP address of the router assigned to be a PIM RP.

The *acl* argument is the name or number of a standard ACL that defines with which multicast groups the RP should be used. 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 bidirectional capability in the PIM hello messages, regardless of the actual bidirectional configuration.

Example:

```
ciscoasa(config)# pim rp-address 10.86.75.23 [acl1] [bidir]
```

Configure 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. You can change this value.

Procedure

Change the designated router priority:

pim dr-priority num

Example:

```
ciscoasa(config-if)# pim dr-priority 500
```

The *num* argument can be any number ranging from 1 to 4294967294.

Configure and Filter 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.

Procedure

Configure the ASA to filter PIM register messages:

pim accept-register {**list** acl | **route-map** map-name}

Example:

```
ciscoasa(config)# pim accept-register {list acl1 | route-map map2}
```

In the example, the ASA filters PIM register messages *acl1* and route map *map2*.

Configure 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.

Procedure

Step 1 Send router query messages:

pim hello-interval seconds

Example:

```
ciscoasa(config-if)# pim hello-interval 60
```

Valid values for the *seconds* argument range from 1 to 3600 seconds.

Step 2 Change the amount of time (in seconds) that the ASA sends PIM join or prune messages:

pim join-prune-interval seconds

Example:

```
ciscoasa(config-if)# pim join-prune-interval 60
```

Valid values for the *seconds* argument range from 10 to 600 seconds.

Filter 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.

Procedure

Step 1 Use a standard ACL to define the routers that you want to have participate in PIM:

access-list pim_nbr deny *router-IP_addr PIM neighbor*

Example:

```
ciscoasa(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255
```

In the example, the following ACL, when used with the **pim neighbor-filter** command, prevents the 10.1.1.1 router from becoming a PIM neighbor.

Step 2 Filter neighbor routers:

pim neighbor-filter pim_nbr

Example:

```
ciscoasa(config)# interface GigabitEthernet0/3
ciscoasa(config-if)# pim neighbor-filter pim_nbr
```

In the example, the 10.1.1.1 router is prevented from becoming a PIM neighbor on interface GigabitEthernet0/3.

Configure 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.

Procedure

Step 1 Use a standard ACL to define the routers that you want to have participate in PIM:

access-list pim_nbr deny *router-IP_addr PIM neighbor*

Example:

```
ciscoasa(config)# access-list pim_nbr deny 10.1.1.1 255.255.255.255
```

In the example, the following ACL, when used with the **pim neighbor-filter** command, prevents the 10.1.1.1 router from becoming a PIM neighbor.

Step 2 Filter neighbor routers:

pim bidirectional-neighbor-filter pim_nbr

Example:

```
ciscoasa(config)# interface GigabitEthernet0/3
ciscoasa(config-if)# pim bidirectional neighbor-filter pim_nbr
```

In the example, the 10.1.1.1 router is prevented from becoming a PIM bidirectional neighbor on interface GigabitEthernet0/3.

Configure 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. 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.

Procedure

Configure a multicast boundary:

multicast boundary acl [**filter-autorp**]

Example:

```
ciscoasa(config-if)# multicast boundary acl1 [filter-autorp]
```

Example for Multicast Routing

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

1. Enable multicast routing:

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

2. Configure a static multicast route:

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

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

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

History for Multicast Routing

Table 1: 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**.

Feature Name	Platform Releases	Feature Information
Protocol Independent Multicast Source-Specific Multicast (PIM-SSM) pass-through support	9.5(1)	Support was added to allow PIM-SSM packets to pass through when multicast routing is enabled, unless the ASA is the Last-Hop Router. This allows greater flexibility in choosing a multicast group while also protecting against different attacks; hosts only receive traffic from explicitly-requested sources. We did not change any commands.