

Multicast

This chapter describes how to configure the Secure Firewall Threat Defense device to use the multicast routing protocol.

- Multicast routing, on page 1
- Prerequisites for Multicast Routing, on page 6
- Guidelines for multicast routing, on page 6
- Configure IGMP features, on page 7
- Configure PIM features, on page 12
- Configure multicast routes, on page 18
- Configure multicast boundary filters, on page 19

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 include videoconferencing, corporate communications, distance learning, and distributing software, stock quotes, and news.

Multicast routing applications and protocols

Multicast routing protocols send source traffic to multiple receivers efficiently, without adding extra load on the source or receivers, and use less network bandwidth than other technologies. Multicast packets are replicated in the network by Firewall Threat Defense device 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 Firewall Threat Defense device supports both stub multicast routing and PIM multicast routing. However, you cannot configure both concurrently on a single Firewall Threat Defense device.

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

IGMP Protocol

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.

Note When you enable multicast routing on the Firewall Threat Defense device, IGMP Version 2 is automatically enabled on all interfaces.

Query Messages to Multicast Groups

The Firewall Threat Defense device 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 Firewall Threat Defense device. If the Firewall Threat Defense device discovers that there are no local members of a multicast group still attached to an interface, it stops forwarding multicast packets 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 Firewall Threat Defense device does not receive a response to a host query within this amount of time, it deletes the group.

Stub multicast routing

Stub multicast routing provides dynamic host registration and facilitates multicast routing.

When configured for stub multicast routing, the Firewall Threat Defense device acts as an IGMP proxy agent. Instead of fully participating in multicast routing, the Firewall Threat Defense device forwards IGMP messages to an upstream multicast router, which sets up delivery of the multicast data. When configured for stub multicast routing, the Firewall Threat Defense device cannot be configured for PIM sparse or bidirectional mode. Enable PIM on the interfaces that participate in IGMP stub multicast routing.

The Firewall Threat Defense device 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, each rooted at a single Rendezvous Point (RP) per multicast group. Optionally, it creates shortest-path trees for each multicast source.

PIM multicast routing

Bidirectional PIM is a variant of PIM-SM that builds bidirectional shared trees connecting multicast sources and receivers.

DF election process

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 then along the shared tree to receivers. This process does not require source-specific state. The DF election takes place during RP discovery and provides a default route to the RP.



Note If the Firewall Threat Defense device is the PIM RP, use the untranslated outside address of the Firewall Threat Defense device as the RP address.

PIM sparse mode

PIM Sparse Mode (PIM-SM) is a PIM mode that forwards traffic only when specifically requested by receivers. It uses a central Rendezvous Point (RP) to coordinate the multicast traffic forwarding between the multicast source and the destination. RP is a router that serves as the meeting point for the source and the receivers. Multicast sources register with the RP through their first hop router. They send a copy of the multicast traffic to the receivers through the RP.

PIM source specific multicast support

Firewall Threat Defense device can forward source specific multicast traffic (e.g., for groups in the 232.x.x.x range) even though they do not support PIM-SSM configuration or Internet Group Management Protocol version 3 (IGMPv3).

SSM is classified as a data delivery mechanism for one-to-many applications such as IPTV. The SSM model uses a concept of "channels" denoted by an (S,G) pair, where S is a source address and G is an SSM destination address. Subscribing to a channel is achieved by using a group management protocol such as IGMPv3. SSM enables a receiving client, once it has learned about a particular multicast source, to receive multicast streams directly from the source rather than receiving it from a shared Rendezvous Point (RP). Access control mechanisms are introduced within SSM, providing a security enhancement that is not available with current sparse or sparse-dense mode implementations.

Limitation:

- Firewall Threat Defense device cannot act as the last-hop router for SSM multicast because it does not support IGMPv3 (which receivers use to join SSM groups).

  If the Firewall Threat Defense device is the last-hop, it will ignore IGMPv3 join messages from receivers for SSM groups, and SSM forwarding will not work.

- Static multicast routes on Firewall Threat Defense device do not work for SSM range (232.x.x.x).

Workaround:

For Cisco Firewall Threat Defense device firewall to forward SSM-related multicast traffic, you must add a multicast-capable layer 3 device (such as a router or switch) that supports PIM and IGMPv3 on the same network segment as the receivers.

How it Works:

- Receivers register their SSM group joins (IGMPv3) with this layer 3 device.
- This layer 3 device sends PIM join messages towards the Firewall Threat Defense device.
- The Firewall Threat Defense device receives these PIM messages and dynamically learns multicast routes.
- SSM multicast traffic is then properly forwarded by the Firewall Threat Defense device, since it is no longer the last-hop device.

Multicast Bidirectional PIM

Multicast bidirectional PIM is useful for networks that have many sources and receivers talking to each other simultaneously and where each participant can become both the source and receiver of multicast traffic, such as in videoconferencing, Webex meetings, and group chat. When PIM bidirectional mode is used, the RP only creates the (*,G) entry for the shared tree. There is no (S,G) entry. This conserves resources on the RP because state tables for each (S,G) entry are not maintained.

In PIM sparse mode, traffic only flows down the shared tree. In PIM bidirectional mode, traffic flows up and down the shared tree.

PIM bidirectional mode also does not use the PIM register/register-stop mechanism to register sources to the RP. Each source can begin sending to the source at any time. When the multicast packets arrive at the RP, they are forwarded down the shared tree (if there are receivers) or dropped (when there are no receivers). However, there is no way for the RP to tell the source to stop sending multicast traffic.

Design-wise you must think about where to place the RP in your network because it should be somewhere in the middle between the sources and receivers in the network.

PIM bidirectional mode has no Reverse Path Forwarding (RPF) check. Instead it uses the concept of a Designated Forwarder (DF) to prevent loops. This DF is the only router on the segment that is allowed to send multicast traffic to the RP. If there is only one router per segment that forwards multicast traffic, there will be no loops. The DF is chosen using the following mechanism:

- The router with the lowest metric to the RP is the DF.
- If the metric is equal, then the router with the highest IP address becomes the DF.

PIM bootstrap router (BSR)

PIM Bootstrap Router (BSR) is a dynamic Rendezvous Point (RP) selection model that uses candidate routers for RP function and for relaying the RP information for a group.

The RP function includes RP discovery and provides a default route to the RP. It does this by configuring a set of devices as candidate BSRs (C-BSR) which participate in a BSR election process to choose a BSR amongst themselves. Once the BSR is chosen, devices that are configured as candidate Rendezvous Points (C-RP) start sending their group mapping to the elected BSR. The BSR then distributes the group-to-RP mapping information to all the other devices down the multicast tree through BSR messages that travel from PIM router to PIM router on a per-hop basis.

This feature allows networks to learn RPs dynamically. Dynamic RP discovery is essential in large, complex networks where an RP may occasionally go down and come up.

PIM bootstrap router (BSR) terminology

This reference provides definitions for key terminology used in PIM Bootstrap Router (BSR) configurations to help understand the roles and functions of different components in the BSR election and RP advertisement process.

These terms are frequently referenced in the PIM BSR configuration:

- Bootstrap Router (BSR) — A BSR advertises Rendezvous Point (RP) information to other routers with PIM on a hop-by-hop basis. Among multiple Candidate-BSRs, a single BSR is chosen after an election process. The primary purpose of this Bootstrap router is to collect all Candidate-RP (C-RP) announcements in to a database called the RP-set and to periodically send this out to all other routers in the network as BSR messages (every 60 seconds).
- Bootstrap Router (BSR) messages — BSR messages are multicast to the All-PIM-Routers group with a TTL of 1. All PIM neighbors that receive these messages retransmit them (again with a TTL of 1) out of all interfaces except the one in which the messages were received. BSR messages contain the RP-set and the IP address of the currently active BSR. This is how C-RPs know where to unicast their C-RP messages.
- Candidate Bootstrap Router (C-BSR) — A device that is configured as a candidate-BSR participates in the BSR election mechanism. A C-BSR with highest priority is elected as the BSR. The highest IP address of the C-BSR is used as a tiebreaker. The BSR election process is preemptive, for example if a new C-BSR with a higher priority comes up, it triggers a new election process.
- Candidate Rendezvous Point (C-RP) — An RP acts as a meeting place for sources and receivers of multicast data. A device that is configured as a C-RP periodically advertises the multicast group mapping information directly to the elected BSR through unicast. These messages contain the Group-range, C-RP address, and a hold time. The IP address of the current BSR is learned from the periodic BSR messages that are received by all routers in the network. In this way, the BSR learns about possible RPs that are currently up and reachable.



Note The Firewall Threat Defense device does not act as a C-RP, even though the C-RP is a mandatory requirement for BSR traffic. Only routers can act as a C-RP. So, for BSR testing functionality, you must add routers to the topology.

- BSR Election Mechanism — Each C-BSR originates Bootstrap messages (BSMs) that contain a BSR Priority field. Routers within the domain flood the BSMs throughout the domain. A C-BSR that hears about a higher-priority C-BSR than itself suppresses its sending of further BSMs for some period of time. The single remaining C-BSR becomes the elected BSR, and its BSMs inform all the other routers in the domain that it is the elected BSR.

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. Hosts can be

located anywhere on the Internet. Hosts that want to receive data sent to a group must join the group using IGMP.

Group membership requirements

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.

Prerequisites for Multicast Routing

Model support

Firewall Threat Defense

Firewall Threat Defense Virtual

Supported domains

Any

User roles

- Admin
- Network Admin

Guidelines for multicast routing

Table 1: System requirements

Item	Description
Context mode	Supported in single context mode.
Firewall mode	Supported only in routed firewall mode. Transparent firewall mode is not supported.

Item	Description
IPv6	IPv6 is not supported.
Multicast group	The range of addresses from 224.0.0.0 to 224.0.0.255 is reserved for routing protocols and other topology discovery or maintenance protocols, such as gateway discovery and group membership reporting. Internet multicast routing from the 224.0.0/24 address range is not supported. When multicast routing is enabled for reserved addresses, an IGMP group is not created.
Clustering	In clustering, for IGMP and PIM, this feature is only supported on the primary unit.

Access control configuration

You must configure an access control or prefilter rule on the inbound security zone to allow traffic to the multicast host, such as 224.1.2.3. However, you cannot specify a destination security zone for the rule, or it cannot be applied to multicast connections during initial connection validation.

Interface and protocol management

Ensure proper configuration management when working with PIM-enabled interfaces and multicast routing protocols.

- You cannot disable an interface with PIM configured on it. If you have configured PIM on the interface (see Configure PIM protocol, on page 13), disabling the multicast routing and PIM does not remove the PIM configuration. You must remove (delete) the PIM configuration to disable the interface.
- Multicast routing is not supported in ECMP.
- Do not configure Firewall Threat Defense to simultaneously be a Rendezvous Point (RP) and a First Hop Router.
- The HSRP standby IP address does not participate in PIM neighborship. If the RP router IP is routed through the HSRP standby IP address, multicast routing does not work in Firewall Threat Defense. To ensure multicast traffic passes through successfully, configure the route for the RP address to use an interface IP address rather than the HSRP standby IP address.
- For a device using virtual routing, you can configure multicast only for its global virtual router and not for its user-defined virtual router.

Configure IGMP features

Configure IGMP features to enable dynamic registration of individual hosts in multicast groups on a particular LAN and to manage multicast traffic efficiently.

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.

IP hosts use IGMP to report their group memberships to directly connected multicast routers. IGMP dynamically registers individual hosts in a multicast group on a specific 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 settings on a per-interface basis.

Procedure

Step 1 Enable multicast routing, on page 8.
Step 2 Configure IGMP protocol, on page 9.
Step 3 Configure IGMP access groups, on page 10.
Step 4 Configure IGMP static groups, on page 11.
Step 5 Configure IGMP join groups, on page 11.

Enable multicast routing

Enable multicast routing to allow the device to forward multicast traffic using Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM) protocols on all interfaces.

Enabling multicast routing on the Firewall Threat Defense device, 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.

These are the maximum number of entries for specific multicast tables. Any new entries are discarded when these limits are reached.

- Multicast Forwarding Information Base (MFIB): 30,000
- IGMP Groups: 30,000
- PIM Routes: 72,000

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.
Step 3 Check the **Enable Multicast Routing** check box.

Checking this check box enables IP multicast routing on the device. Unchecking it disables IP multicast routing. By default, multicast is disabled. Enabling multicast routing enables multicast on all interfaces.

You can disable multicast on a per-interface basis. This is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the Firewall Threat Defense device from sending host query messages on that interface.

Configure IGMP protocol

Use this task to configures IGMP parameters per-interface. These settings help manage multicast group membership on a Firewall Threat Defense device.

You can configure IGMP parameters per interface, such as the forward interface, query messages, and time intervals.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 Click the **Protocol** tab and click **Add**, or edit an interface.

In the **Add IGMP parameters** dialog box, configure new IGMP parameters to the Firewall Threat Defense device. Use the **Edit IGMP parameters** dialog box to modify the existing parameters.

Step 4 Configure these parameters:

- **Interface**: From the drop-down list, choose the interface for which you want to configure IGMP protocol.
- **Enable IGMP**: Check the check box to enable IGMP.

  Note
  Disabling IGMP on specific interfaces is useful if you know that there are no multicast hosts on a specific interface and you want to prevent the device from sending host query messages on that interface.

- **Forward Interface**: From the drop-down list, choose the specific interface from which you want to forward IGMP messages.

  This configures the Secure Firewall Threat Defense device 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.

- **Version**: Choose IGMP Version 1 or 2.

  By default, the Firewall Threat Defense device runs IGMP Version 2, which enables several additional features.

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

- **Query Interval**: The interval in seconds at which the designated router sends IGMP host-query messages. The range is 1 to 3600. The default is 125.

Note
If the Firewall Threat Defense device does not hear a query message on an interface for the specified timeout value, then the device becomes the designated router and starts sending the query messages.

- **Response Time**: The interval in seconds before the Firewall Threat Defense device deletes the group. The range is 1 to 25. The default is 10.

  If the Firewall Threat Defense device does not receive a response to a host query within this amount of time, it deletes the group.

- **Group Limit**: Maximum number of hosts that can join on an interface. The range is 1 to 500. The default is 500.

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

- **Query Timeout**: The period of time in seconds before which the Firewall Threat Defense device takes over as the requester for the interface after the previous requester has stopped. The range is 60 to 300. The default is 255.

Step 5 Click **OK** to save the IGMP protocol configuration.

Configure IGMP access groups

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

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Access Group**.

Step 3 On **Access Group**, click **Add** or **Edit**.

Use the **Add IGMP Access Group parameters** dialog box to add new IGMP access groups to the Access Group table. Use the **Edit IGMP Access Group parameters** dialog box to change the parameters of an existing access group.

Step 4 Configure the access group options:

a) From the **Interface** drop-down list, choose the interface associated with the access group. You cannot change the associated interface when you are editing an existing access group.

b) Click one of the following:

- **Standard Access List**: From the **Standard Access List** drop-down list, choose the standard ACL or click **Add** (+) to create a new standard ACL. For more information, refer to Configure a standard ACL object.
- **Extended Access List**: From the **Extended Access List** drop-down list, choose the extended ACL or click **Add** (+) to create a new extended ACL. For more information, refer to Configure an extended ACL object.

Step 5 Click **OK** to save the access group configuration.

Configure IGMP static groups

Configure a statically joined IGMP group to enable multicast traffic for that group to be sent to a network segment even when group members cannot report membership or no members exist on the segment.

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 by configuring a statically joined IGMP group. With this method, the Firewall Threat Defense device does not accept the packets itself, but only forwards them. Therefore, this method allows fast switching. The outgoing interface appears in the IGMP cache, but this interface is not a member of the multicast group.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Static Group**, click **Add** or **Edit**.

Use the **Add IGMP Static Group parameters** dialog box to statically assign a multicast group to an interface. Use the **Edit IGMP Static Group parameters** dialog box to change existing static group assignments.

Note
The IGMP Static Group enables PIM to send *Join* requests towards the sources or towards the Rendezvous Point (RP), provided, the firewall with this command is the PIM Designated Router (DR) on that interface where the command is applied.

Step 4 Configure the following options:

- From the **Interface** drop-down list, choose the interface to which you want to statically assign a multicast group. If you are editing an existing entry, you cannot change the value.
- From the **Multicast Groups** drop-down list, choose the multicast group to which you want to assign the interface, or click **Add** (+) to create a new multicast group. Refer to Creating Network Objects for the procedure.

Step 5 Click **OK** to save the static group configuration.

Configure IGMP join groups

You can configure an interface to be a member of a multicast group.

Configuring the Firewall Threat Defense device to join a multicast group, upstream routers maintain multicast routing table information for that group and keep the group's paths active.

Note Refer to Configure IGMP static groups, on page 11 if you want to forward multicast packets for a specific group to an interface without the Firewall Threat Defense device accepting those packets as part of the group.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **IGMP**.

Step 3 On **Join Group**, click **Add** or **Edit**.

Use the **Add IGMP Join Group parameters** dialog box to configure the Firewall Threat Defense device to be a member of a multicast group. Use the **Edit IGMP Join Group parameters** dialog box to change existing parameters.

Note
The IGMP Join Group enables PIM to send *Join* requests to sources or to the Rendezvous Point (RP) if the firewall with this command is the PIM Designated Router (DR) on that interface.

Step 4 Configure the following options:

- From the **Interface** drop-down list, choose the interface you want to be a member of a multicast group. If you are editing an existing entry, you cannot change the value.
- From the **Join Group** drop-down list, choose the multicast group to which you want to assign the interface, or click **Plus** to create a new multicast group. Refer to Creating Network Objects for the procedure.

Configure PIM features

This section describes how to configure optional PIM settings.

Routers use PIM to maintain forwarding tables to use for forwarding multicast diagrams. Enable multicast routing on the Secure Firewall Threat Defense device, to automatically enable IGMP and PIM 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.

Procedure

Step 1 Configure PIM protocol, on page 13.

Step 2 Configure PIM neighbor filters, on page 14.

Step 3 Configure PIM bidirectional neighbor filters, on page 14.

Step 4 Configure PIM rendezvous points, on page 15.

Step 5 Configure PIM route trees, on page 16.

Step 6 Configure PIM request filters, on page 17.

Step 7 Configure multicast boundary filters, on page 19.

Configure PIM protocol

You can enable or disable PIM on a specific interface.

You can also configure the Designated Router (DR) 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, the DR is chosen based on DR priority. If multiple devices share the same DR priority, the device with the highest IP address becomes the DR. By default, the Firewall Threat Defense device has a DR priority of 1.

Router query messages are used to choose 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 Firewall Threat Defense device sends PIM join or prune messages.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 Click the **Protocol** tab and click **Add**, or edit an interface.

In the **Add PIM parameters** dialog box, add new PIM parameters for the interface. Use the **Edit PIM parameters** dialog box to change existing parameters.

Step 4 Configure these options:

- **Interface**: From the drop-down list, select the interface for which you want to configure PIM protocol.
- **Enable PIM**: Check the check box to enable PIM.
- **DR Priority**: Set the DR priority value for the selected interface. The router with the highest DR priority on the subnet becomes the designated router.
- **Hello Interval**: The interval in seconds at which the interface sends PIM hello messages. The range is 1 to 3600. The default is 30.
- **Join Prune Interval**: The interval in seconds at which the interface sends PIM join and prune advertisements. The range is 10 to 600. The default is 60.

Step 5 Click **OK** to save the PIM protocol configuration.

Configure PIM neighbor filters

You can define and filter the routers that become Protocol Independent Multicast (PIM) neighbors.

PIM neighbor filters allow you to define and filter the routers that become Protocol Independent Multicast (PIM) neighbors to:

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

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 Click the **Neighbor Filter** tab and click **Add**, or edit an interface.

In the **Add PIM Neighbor Filter** dialog box, add new PIM neighbor filters to the interface. Use the **Edit PIM Neighbor Filter** dialog box to change existing parameters.

Step 4 Configure these options:

- From the **Interface** drop-down list, choose the interface to which you want to add a PIM neighbor filter.
- From the **Standard Access List** drop-down list, choose a standard ACL or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object for the procedure.

  Note
  Choosing **Allow** on the **Add Standard Access List Entry** dialog box lets the multicast group advertisements pass through the interface. Choosing **Block** prevents the specified multicast group advertisements from passing through the interface. When a multicast boundary is configured on an interface, all multicast traffic is prevented from passing through the interface unless permitted with a neighbor filter entry.

Step 5 Click **OK** to save the PIM neighbor filter configuration.

Configure PIM bidirectional neighbor filters

Configure PIM bidirectional neighbor filters to control which neighbor devices participate in the Designated Forwarder (DF) election for bidirectional multicast routing. This helps ensure proper multicast traffic forwarding in the network.

A PIM bidirectional neighbor filter is an ACL that defines the neighbor devices that can participate in the Designated Forwarder (DF) election. If you do not configure a PIM bidirectional neighbor filter for an interface, there are no restrictions. If you configure a PIM bidirectional neighbor filter, only neighbors permitted by the ACL can participate in the DF election process.

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

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

- If a permitted neighbor does not support bidirectional mode, the DF election does not occur.
- If a denied neighbor supports bidirectional mode, the DF election does not occur.
- If a denied neighbor does not support bidirectional mode, the DF election can occur.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Multicast Routing** > **PIM**.

Step 3 On **Bidirectional Neighbor Filter**, click **Add** or **Edit**.

Use the **Add PIM Bidirectional Neighbor Filter** dialog box to create ACL entries for the PIM bidirectional neighbor filter ACL. Use the **Edit PIM Bidirectional Neighbor Filter** dialog box to change existing parameters.

Step 4 Configure these options:

- In the **Interface** drop-down list, select the interface for which you want to configure the PIM bidirectional neighbor filter ACL entry.
- From the **Standard Access List** drop-down list, select a standard ACL or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object for the procedure.

  Note
  Choosing **Allow** in the **Add Standard Access List Entry** dialog box lets the specified devices participate in the DR election process. Choosing **Block** prevents the specified devices from participating in the DR election process.

Step 5 Click **OK** to save the PIM bidirectional neighbor filter configuration.

Configure PIM rendezvous points

You can configure the Firewall Threat Defense device to serve as a 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). Refer to Multicast Bidirectional PIM, on page 4 for more information about bidirectional PIM.

These restrictions apply to RPs:

- You cannot use the same RP address twice.
- You cannot specify **All Groups** for more than one RP.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Rendezvous Points**, click **Add** or **Edit**.

Use the **Add Rendezvous Point** dialog box to create a new entry in the Rendezvous Point table. Use the **Edit Rendezvous Point** dialog box to modify existing parameters.

Step 4 Configure these options:

- From the **Rendezvous Point IP address** drop-down list, choose the IP address that you want to add as an RP or click **Add** (+) to create a new network object. Refer to Creating Network Objects for the procedure.
- Check the **Use bi-directional forwarding** check box to enable bidirectional mode for the specified multicast groups. In this mode, if the Firewall Threat Defense device receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a prune message back to the source.
- Click **Use this RP for all Multicast Groups** to use the specified RP for all multicast groups on the interface.
- Click **Use this RP for all Multicast Groups as specified below** to designate the multicast groups to use with the specified RP and then from the **Standard Access List** drop-down list, choose a standard ACL or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object for the procedure.

Step 5 Click **OK** to save the rendezvous point configuration.

Configure PIM route trees

Configure PIM route trees to control whether the Firewall Threat Defense device joins the shortest-path tree or uses the shared tree for multicast traffic routing.

By default, PIM leaf routers join the shortest-path tree immediately after the first packet arrives from a new source. This method reduces delay, but requires more memory than the shared tree. You can configure whether or not the Firewall Threat Defense device should join the shortest-path tree or use the shared tree, either for all multicast groups or only for specific multicast addresses.

The shortest-path tree is used for any group that is not specified in the Multicast Groups table. The Multicast Groups table displays the multicast groups to use with the shared tree. The table entries are processed from the top down. You can create an entry that includes a range of multicast groups, but excludes specific groups within that range by placing deny rules for the specific groups at the top of the table and the permit rule for the range of multicast groups below the deny statements.

Note This behavior is known as Shortest Path Switchover (SPT). We recommend that you always use the Shared Tree option.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Route Tree**, select the path for the route tree:

- Click **Shortest Path** to use the shortest-path tree for all multicast groups.
- Click **Shared Tree** to use the shared tree for all multicast groups.
- Click **Shared tree for below mentioned group** to designate the groups specified in the Multicast Groups table, and then from the **Standard Access List** drop-down list, select a standard ACL or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object for the procedure.

Step 4 Click **OK** to save the route tree configuration.

Configure PIM request filters

When the Firewall Threat Defense device is acting as an RP, you can restrict specific multicast sources from registering with it to prevent unauthorized sources from registering with the RP. You can define the multicast sources from which the Firewall Threat Defense device will accept PIM register messages.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Request Filter**, define the multicast sources that are allowed to register with the Firewall Threat Defense device when it acts as an RP:

- From the **Filter PIM register messages using:** drop-down list select **None**, **Access List**, or **Route Map**.
- If you choose **Access List** from the drop-down list, select an extended ACL or click **Add** (+) to create a new extended ACL. Refer to Configure an extended ACL object for the procedure.

  Note
  In the **Add Extended Access List Entry** dialog box, select **Allow** from the drop-down list to create a rule that allows the specified source of the specified multicast traffic to register with the Firewall Threat Defense device, or select **Block** to create a rule that prevents the specified source of the specified multicast traffic from registering with the device.

- If you choose **Route Map**, select a route map from the **Route Map** drop-down list, or click **Add** (+) to create a new route map. Refer to Creating Network Objects for the procedure.

Step 4 Click **OK** to save the request filter configuration.

Configure the Firewall Threat Defense device as a candidate bootstrap router

You can configure the Firewall Threat Defense device as a candidate BSR.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **PIM**.

Step 3 On **Bootstrap Router**, check the **Configure this FTD as a Candidate Bootstrap Router (C-BSR)** check box to perform the C-BSR setup.

a) From the **Interface** drop-down list, select the interface in the Firewall Threat Defense device from which the BSR address is derived to make it a candidate.

This interface must be enabled with PIM.

b) In the **Hash mask length** field, enter the length of a mask (32 bits maximum) that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash (correspond) to the same RP. For example, if this value is 24, only the first 24 bits of the group addresses matter. This allows you to configure one RP for multiple groups. The range is 0 to 32.

c) In the **Priority** field, enter the priority of the candidate BSR. The BSR with the larger priority is preferred. If the priority values are the same, the router with the larger IP address is the BSR. The range is 0 to 255. The default value is 0.

Step 4 (Optional) Click **Add** (+) to select an interface on which no PIM BSR messages will be sent or received in the **Configure this FTD as a Border Bootstrap Router (BSR)** section.

- From the **Interface** drop-down list, select the interface on which no PIM BSR messages will be sent or received.

  RP or BSR advertisements are filtered effectively isolating two domains of RP information exchange.

- Check the **Enable Border BSR** check box to enable BSR.

Step 5 Click **OK** to save the bootstrap router configuration.

Configure multicast routes

Configure static multicast routes to 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 Firewall Threat Defense device 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 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Multicast Routes** > **Add or Edit**.

Use the **Add Multicast Route Configuration** dialog box to add a new multicast route to the Firewall Threat Defense device. Use the **Edit Multicast Route Configuration** dialog box to change an existing multicast route.

Step 3 Choose an existing network from the **Source Network** drop-down, or click **Add** (+) to add a new network.

Refer to Creating Network Objects for the procedure.

Step 4 To configure an interface to forward the route, click **Interface** and configure the following options:

- From the **Source Interface** drop-down list, choose the incoming interface for the multicast route.
- From the **Output Interface/Dense** drop-down list, choose the destination interface that the route is forwarded through.
- In the **Distance** field, enter the distance of the multicast route. The range is 0 to 255.

Step 5 To configure an RPF address to forward the route, click **Address** and configure the following options:

- In the **RPF Address** field, enter the IP address for the multicast route.
- Enter the multicast route distance in the **Distance** field. The value must be between 0 and 255.

Step 6 Click **OK** to save the multicast routes configuration.

Configure multicast boundary filters

Address scoping defines domain boundary filters 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 filter 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 are considered local and are not globally unique.

A standard ACL defines the range of affected addresses. When a boundary filter is set up, no multicast data packets are allowed to flow across the boundary from either direction. The boundary filter 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. 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 filter 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

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Choose **Routing** > **Multicast Routing** > **Multicast Boundary Filter**, and then click **Add** or **Edit**.

Use the **Add Multicast Boundary Filter** dialog box to add new multicast boundary filters to the device. Use the **Edit Multicast Boundary Filter** dialog box to change existing parameters.

You can configure a multicast boundary for administratively scoped multicast addresses. A multicast boundary restricts multicast data packet flows and enables reuse of the same multicast group address in different administrative domains. When a multicast boundary is defined on an interface, only the multicast traffic permitted by the filter ACL passes through the interface.

Step 3 From the **Interface** drop-down list, choose the interface for which you are configuring the multicast boundary filter ACL.

Step 4 From the **Standard Access List** drop-down list, choose the standard ACL you want to use, or click **Add** (+) to create a new standard ACL. Refer to Configure a standard ACL object for the procedure.

Step 5 Check the **Remove any Auto-RP group range announcement from the Auto-RP packets that are denied by the boundary** check box to filter Auto-RP messages from sources denied by the boundary ACL. If this check box is not checked, all Auto-RP messages are passed.

Step 6 Click **OK** to save the multicast boundary filter configuration.

The multicast boundary filter is configured on the selected interface with the specified standard ACL, restricting multicast traffic according to the filter rules and enabling address reuse across administrative domains.