



Configuring Routing Policies on Firewall Devices

The Routing section in Security Manager contains pages for defining and managing routing settings for security appliances.

This chapter contains the following topics:

- [Configuring No Proxy ARP](#) , on page 1
- [Configuring BGP](#) , on page 2
- [Configuring EIGRP](#) , on page 34
- [Configuring ISIS](#) , on page 50
- [Configuring BFD Routing](#) , on page 72
- [Configuring OSPF](#) , on page 80
- [Configuring Key Chain](#) , on page 108
- [Configuring OSPFv3](#) , on page 112
- [Configuring RIP](#) , on page 131
- [Configuring Static Routes](#) , on page 141
- [Configuring Policy Objects for ASA Routing Policies](#) , on page 144

Configuring No Proxy ARP

When a host sends IP traffic to another device on the same Ethernet network, the host needs to know the MAC address of the device. Address Resolution Protocol (ARP) is a Layer 2 protocol that resolves an IP address to a MAC address: a host sends an ARP request asking “Who is this IP address?” The device owning the IP address replies, “I own that IP address; here is my MAC address.”

With Proxy ARP, a device responds to an ARP request with its own MAC address, even though the device does not own the IP address. Serving as an ARP Proxy for another host effectively directs network traffic to the proxy, in this case your security appliance. Traffic that passes through the appliance is then routed to the appropriate destination.

For example, the security appliance uses proxy ARP when you configure NAT and specify a global address that is on the same network as the appliance interface. The only way traffic can reach the destination hosts is if the appliance claims and subsequently routes traffic to the destination global addresses.

By default, proxy ARP is enabled for all interfaces. Use the No Proxy ARP page to disable proxy ARP for global addresses:

- To disable proxy ARP for one or more interfaces, enter their names in the Interfaces field. Separate multiple interfaces with commas. You can click Select to choose the interfaces from a list of interfaces defined on the device, and interface roles defined in Security Manager.



Note On ASA 8.4.2 and later devices operating in routed mode, you can disable Proxy ARP on the egress interface for a Manual NAT rule. See *Do not proxy ARP on Destination Interface* in Table 24-15 for more information.

Navigation Path

- (Device view) Select **Platform > Routing > No Proxy ARP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > No Proxy ARP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- [Configuring Static Routes](#) , on page 141
- [Configuring RIP](#) , on page 131
- [Configuring OSPF](#) , on page 80

Configuring BGP

Border Gateway Protocol (BGP) is an inter autonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).



Note BGP configuration is supported on ASA 9.2(1)+ only. Also, beginning with ASA 9.3(1), BGP is supported in L2 (EtherChannel Type) and L3 (Individual Interface Type) clustering modes.

Navigation Path

- (Device view) Select **Platform > Routing > BGP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > BGP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

The BGP page provides two tabbed panels for configuring BGP routing on a firewall device. This is the basic procedure for configuring the BGP process:

1. Enable the BGP routing process by checking the Enable BGP check box on the BGP page.
2. In the AS Number field, enter the autonomous system (AS) number for the BGP process. The AS number internally includes multiple autonomous numbers. The AS number can be from 1 to 4294967295 or from 1.0 to 65535.65535.

3. On the [General Tab](#) , on page 5:
 - (Optional) Check the Limit the number of AS numbers in AS_PATH attribute of received routes check box to restrict the number of AS numbers in AS_PATH attribute to a specific number. Valid values are from 1 to 254.
 - (Optional) Check the Log Neighbor Changes check box to enable logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability.
 - (Optional) Check the Use TCP path MTU Discovery check box to use the Path MTU Discovery technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation.
 - (Optional) Check the Enable fast external failover check box to reset the external BGP session immediately upon link failure.
 - (Optional) Check the Enforce that the first AS is peer's AS for EBGP routes check box to discard incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system.
 - (Optional) Check the Use dot notation for AS numbers check box to split the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65553 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation.
 - Define the configuration related to the best path selection process for BGP routing (see [General Tab](#) , on page 5).
 - Specify the timer information in the Neighbor timers area (see [General Tab](#) , on page 5).
 - (Optional) Configure Graceful Restart (see [General Tab](#) , on page 5).
4. On the IPv4 Family tab, select the Enable IPv4 Family check box and then use the tabs provided to configure IPv4 Address Family settings. For more information, see [IPv4 Family Tab](#) , on page 7.
5. On the IPv6 Family tab, select the Enable IPv6 Family check box and then use the tabs provided to configure IPv6 Address Family settings. For more information, see [IPv6 Family Tab](#) , on page 21.

Related Topics

- [About BGP](#) , on page 3

About BGP

BGP is an inter autonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).

When to Use BGP

Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP) such as OSPF for the exchange of routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS),

the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

Routing Table Changes

BGP neighbors exchange full routing information when the TCP connection between neighbors is first established. When changes to the routing table are detected, the BGP routers send to their neighbors only those routes that have changed. BGP routers do not send periodic routing updates, and BGP routing updates advertise only the optimal path to a destination network.

Routes learned via BGP have properties that are used to determine the best route to a destination, when multiple paths exist to a particular destination. These properties are referred to as BGP attributes and are used in the route selection process:

- **Weight** -- This is a Cisco-defined attribute that is local to a router. The weight attribute is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight is preferred.
- **Local preference** -- The local preference attribute is used to select an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS. If there are multiple exit points from the AS, the exit point with the highest local preference attribute is used as an exit point for a specific route.
- **Multi-exit discriminator** -- The multi-exit discriminator (MED) or metric attribute is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. It is referred to as a suggestion because the external AS that is receiving the MEDs may also be using other BGP attributes for route selection. The route with the lower MED metric is preferred.
- **Origin** -- The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three possible values and is used in route selection.
 - **IGP**- The route is interior to the originating AS. This value is set when the network router configuration command is used to inject the route into BGP.
 - **EGP**-The route is learned via the Exterior Border Gateway Protocol (EBGP).
 - **Incomplete**- The origin of the route is unknown or learned in some other way. An origin of incomplete occurs when a route is redistributed into BGP.
- **AS_path** -- When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. Only the route with the shortest AS_path list is installed in the IP routing table.
- **Next hop** -- The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers. For IBGP, the EBGP next-hop address is carried into the local AS.
- **Community** -- The community attribute provides a way of grouping destinations, called communities, to which routing decisions (such as acceptance, preference, and redistribution) can be applied. Route maps are used to set the community attribute. The predefined community attributes are as follows:
 - **no-export**- Do not advertise this route to EBGP peers.
 - **no-advertise**- Do not advertise this route to any peer.
 - **internet**- Advertise this route to the Internet community; all routers in the network belong to it.

BGP Path Selection

BGP may receive multiple advertisements for the same route from different sources. BGP selects only one path as the best path. When this path is selected, BGP puts the selected path in the IP routing table and propagates the path to its neighbors. BGP uses the following criteria, in the order presented, to select a path for a destination:

- If the path specifies a next hop that is inaccessible, drop the update.
- Prefer the path with the largest weight.
- If the weights are the same, prefer the path with the largest local preference.
- If the local preferences are the same, prefer the path that was originated by BGP running on this router.
- If no route was originated, prefer the route that has the shortest AS_path.
- If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
- If the origin codes are the same, prefer the path with the lowest MED attribute.
- If the paths have the same MED, prefer the external path over the internal path.
- If the paths are still the same, prefer the path through the closest IGP neighbor.
- If both paths are external, prefer the path that was received first (the oldest one).
- Prefer the path with the lowest IP address, as specified by the BGP router ID.
- If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length.
- Prefer the path that comes from the lowest neighbor address.

General Tab

Use the General tab to configure BGP settings such as Best Path Selection, Neighbor Timers, and Graceful Restart.

Navigation Path

You can access the Neighbors tab from the BGP page (see [Configuring BGP](#) , on page 2).

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv4 Family Tab](#) , on page 7

Field Reference

Table 1: General Tab

Element	Description
Limit the number of AS numbers in AS_PATH attribute of received routes	Restricts the number of AS numbers in AS_PATH attribute to a specific number. Valid values are from 1 to 254.
Log Neighbor Changes	Enables logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability.
Use TCP path MTU Discovery	Enables the use of the Path MTU Discovery technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation.
Enable fast external failover	Resets the external BGP session immediately upon link failure.
Enforce that the first AS is peer's AS for EBGp routes	Discards incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system.
Use dot notation for AS numbers	Splits the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65533 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation.
Best Path Selection	
Default local preference	Specify a value between 0 and 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers and access servers in the local autonomous system.
Allow comparing MED from different neighbors	Allows the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
Compare Router-id for identical EBGp paths	Compares similar paths received from external BGP peers during the best path selection process and switches the best path to the route with the lowest router ID.
Pick the best MED path among paths advertised from the neighboring AS	Enables MED comparison among paths learned from confederation peers. The comparison between MEDs is made only if no external autonomous systems are there in the path.
Treat missing MED as the least preferred one	Considers the missing MED attribute as having a value of infinity, making the path the least desirable; therefore, a path with a missing MED is least preferred.
Neighbor Timers	

Element	Description
Keepalive Interval	Enter the time interval for which the BGP neighbor remains active after not sending a keepalive message. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.
Hold Time	Enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. The default values is 180 seconds.
Min Hold Time	(Optional) Enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. Specify a value from 0 to 65535.
Graceful Restart (Use in failover or spanned cluster mode) (ASA 9.3.1+ only)	
Enable Graceful Restart	Enables ASA peers to avoid a routing flap following a switchover.
Restart Time	Specify the time duration that ASA peers will wait to delete stale routes before a BGP open message is received. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
Stalepath Time	Enter the time duration that the ASA will wait before deleting stale routes after an end of record (EOR) message is received from the restarting ASA. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.

IPv4 Family Tab

Use the IPv4 Family tab on the BGP page to enable and configure IPv4 settings for BGP.

Navigation Path

You can access the IPv4 Family tab from the BGP page. For more information about the BGP page, see [Configuring BGP](#), on page 2.

Related Topics

- [About BGP](#), on page 3
- [General Tab](#), on page 5

Field Reference

Table 2: IPv4 Family - Aggregate Address Tab

Element	Description
Enable IPv4 Family	Enables configuration of routing sessions that use standard IPv4 address prefixes.

Element	Description
General	Use this panel to configure general IPv4 settings such as Best Path Selection, Neighbor Timers, and Graceful Restart. See IPv4 Family - General Tab , on page 8 for more about these definitions.
Aggregate Address	Use this panel to define the aggregation of specific routes into one route. Specify a value for the aggregate timer (in seconds) in the Aggregate Timer field. Valid values are 0 or any value between 6 and 60. The default value is 30. See Add/Edit Aggregate Address Dialog Box , on page 24 for more about these definitions.
Filtering	Use this panel to filter routes or networks received in incoming BGP updates. See Add/Edit Filter Dialog Box , on page 11 for more about these definitions.
Neighbor	Use this panel to define BGP neighbors and neighbor settings. See Add/Edit Neighbor Dialog Box , on page 12 for more about these definitions.
Networks	Use this panel to define the networks to be advertised by the BGP routing process. See Add/Edit Network Dialog Box , on page 18 for more about these definitions.
Redistribution	Use this panel to define the conditions for redistributing routes from another routing domain into BGP. See Add/Edit Redistribution Dialog Box , on page 19 for more about these definitions.
Route Injection	Use this panel to define the routes to be conditionally injected into the BGP routing table. See Add/Edit Route Injection Dialog Box , on page 20 for more about these definitions.

IPv4 Family - General Tab

Use the IPv4 Family - General tab to configure the general IPv4 settings.

Navigation Path

You can access the General tab from the IPv4 Family Tab on the BGP page. For more information about the IPv4 Family tab, see [IPv4 Family Tab](#) , on page 7.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3

Field Reference

Table 3: IPv4 Family - General Tab

Element	Description
Router ID	<p>On a single device, choose Automatic or IP Address. (An address field appears when you choose IP Address.)</p> <p>If you choose Automatic, the highest-level IP address on the security appliance is used as the router ID. To use a fixed router ID, choose IP Address and enter an IPv4 address in the Router ID field.</p> <p>On a device cluster, choose Automatic or Cluster Pool. (An IPv4 Pool object ID field appears when you choose Cluster Pool.)</p> <p>If you choose Cluster Pool, enter or Select the name of the IPv4 Pool object that is to supply the Router ID address. For more information, see Add or Edit IPv4 Pool Dialog Box.</p>
Learned Route Map	<p>Enter or Select the name of a route map object.</p> <p>Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, on page 145.</p>
Scanning Interval	Enter a scanning interval (in seconds) for BGP routers for next-hop validation. Valid values are from 5 to 60 seconds. The default value is 60.
Routes and Synchronization	
Generate Default Route	(Optional) Configures a BGP routing process to distribute a default route (network 0.0.0.0).
Summarize subnet routes into network-level routes	(Optional) Configures automatic summarization of subnet routes into network-level routes.
Advertise inactive routes	(Optional) Advertises routes that are not installed in the routing information base (RIB).
Synchronize between BGP and the Interior Gateway Protocol (IGP) system	<p>Enables synchronization between BGP and your Interior Gateway Protocol (IGP) system. To enable the Cisco IOS software to advertise a network route without waiting for the IGP, deselect this option.</p> <p>Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. By default, synchronization between BGP and the IGP is turned off to allow the Cisco IOS software to advertise a network route without waiting for route validation from the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems. Use synchronization if routers in the autonomous system do not speak BGP.</p>
Redistribute iBGP into an IGP	(Optional) Configures iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.

Element	Description
Administrative Route Distances	
External	Specifies the administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255. The default value is 20.
Internal	Specifies administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255. The default value is 200.
Local	Specifies administrative distance for local BGP routes. Local routes are those networks listed with a network router configuration command, often as back doors, for the router or for the networks that is being redistributed from another process. The range of values for this argument are from 1 to 255. The default value is 200.
Next Hop	
Enable address tracking	(Optional) Enables BGP next hop address tracking.
Delay Interval	Specify the delay interval between checks on updated next-hop routes installed in the routing table.
Forward packets over Multiple Paths	
Number of Paths	(Optional) Specify the maximum number of external BGP routes that can be installed to the routing table.
IBGP Number of Paths	(Optional) Specify the maximum number of internal BGP routes that can be installed to the routing table.

Add/Edit Aggregate Address Dialog Box

Use the Add/Edit Aggregate Address dialog box to define the aggregation of specific routes into one route.

Navigation Path

You can access the Add/Edit Aggregate Address dialog box from the [IPv4 Family Tab](#) , on page 7.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv4 Family - General Tab](#) , on page 8
- [Add/Edit Filter Dialog Box](#) , on page 11
- [Add/Edit Neighbor Dialog Box](#) , on page 12
- [Add/Edit Network Dialog Box](#) , on page 18
- [Add/Edit Redistribution Dialog Box](#) , on page 19

- [Add/Edit Route Injection Dialog Box](#) , on page 20

Field Reference

Table 4: Add/Edit Aggregate Address Dialog Box

Element	Description
Network	Enter an IP address, or enter or Select the desired Network/Hosts objects.
Attribute Map	(Optional) Enter or Select the route map used to set the attribute of the aggregate route. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Advertise Map	(Optional) Enter or Select the route map used to select the routes to create AS_SET origin communities. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Suppress Map	(Optional) Enter or Select the route map used to select the routes to be suppressed. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Generate AS Set Path Information	Enables generation of autonomous system set path information.
Filter all more-specific routes from updates	Filters all more-specific routes from updates.

Add/Edit Filter Dialog Box

Use the Add/Edit Filter dialog box to filter routes or networks received in incoming BGP updates.

Navigation Path

You can access the Add/Edit Filter dialog box from the [IPv4 Family Tab](#) , on page 7.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv4 Family Tab—General Tab](#), on page 54

- [Add/Edit Aggregate Address Dialog Box](#) , on page 24
- [Add/Edit Neighbor Dialog Box](#) , on page 25
- [Add/Edit Network Dialog Box](#) , on page 31
- [Add/Edit Redistribution Dialog Box](#) , on page 32
- [Add/Edit Route Injection Dialog Box](#) , on page 33

Field Reference

Table 5: Add/Edit Filter Dialog Box

Element	Description
ACL	Select an Access Control List that defines which networks are to be received and which are to be suppressed in routing updates.
Direction	Choose a direction from the Direction drop-down list. The direction will specify if the filter should be applied to inbound updates or outbound updates.
Protocol	Select the routing process for which you want to filter: None, BGP, Connected, EIGRP, OSPF, RIP, or Static.
AS Number	Shows the autonomous system number of the BGP routing process. This value is specified on the BGP page (see Configuring BGP , on page 2).
Process ID	Enter the identifier for the routing process. Applies to EIGRP and OSPF routing protocols.

Add/Edit Neighbor Dialog Box

Use the Add/Edit Neighbor dialog box to define BGP neighbors and neighbor settings.

Navigation Path

You can access the Add/Edit Neighbor dialog box from the [IPv4 Family Tab](#) , on page 7.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv4 Family - General Tab](#) , on page 8
- [Add/Edit Aggregate Address Dialog Box](#) , on page 10
- [Add/Edit Network Dialog Box](#) , on page 18
- [Add/Edit Redistribution Dialog Box](#) , on page 19
- [Add/Edit Route Injection Dialog Box](#) , on page 20

Field Reference

Table 6: Add/Edit Neighbor Dialog Box

Element	Description
General	
IP Address	Enter the BGP neighbor IP address. This IP address is added to the BGP neighbor table.
Remote AS	Enter the autonomous system to which the BGP neighbor belongs.
Enable Address Family	(Optional) Enables communication with the BGP neighbor.
Shutdown neighbor administratively	(Optional) Disable a neighbor or peer group.
Configure Graceful Restart per neighbor (ASA 9.3.1+ only)	(Optional) Enables configuration of the Border Gateway Protocol (BGP) graceful restart capability for this neighbor. After selecting this option, you must use the Graceful Restart (Use in failover or spanned cluster mode) option to specify whether graceful restart should be enabled or disabled for this neighbor.
Graceful Restart (Use in failover or spanned cluster mode) (ASA 9.3.1+ only)	(Optional) Enables the Border Gateway Protocol (BGP) graceful restart capability for this neighbor.
Description	(Optional) Enter a description for the BGP neighbor.
fall-over BFD	(Optional) Enables BFD support for fall-over for the BGP neighbor.
BFD-Hop	(Optional) Specify if there is a single IP hop or multiple IP hops between a BFD source and destination.
Filtering	
Filter routes using an access list	(Optional) Enter or Select the appropriate incoming or outgoing access control list to distribute BGP neighbor information.
Filter routes using route map	(Optional) Enter or Select the appropriate incoming or outgoing route maps to apply a route map to incoming or outgoing routes. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Filter routes using a Prefix list	(Optional) Enter or Select the appropriate incoming or outgoing prefix list to distribute BGP neighbor information. Tip Click Select to open the Prefix List Object Selector from which you can select a prefix list object. You can also create new objects from the object Prefix List Object selector. For more information, see Add or Edit Prefix List Object Dialog Box , on page 159.

Element	Description
Filter routes using AS Path filter	<p>(Optional) Enter or Select the appropriate incoming or outgoing AS path filter to distribute BGP neighbor information.</p> <p>Tip Click Select to open the AS Path Object Selector from which you can select an AS path object. You can also create new AS path objects from the AS Path Object Selector. For more information, see Add or Edit As Path Object Dialog Boxes , on page 163.</p>
Limit the number of prefixes allowed from the neighbor	<p>(Optional) Select to control the number of prefixes that can be received from a neighbor.</p> <ul style="list-style-type: none"> • Enter the maximum number of prefixes allowed from a specific neighbor in the Maximum Prefixes field. • Enter the percentage (of maximum) at which the router starts to generate a warning message in the Threshold Level field. Valid values are integers between 1 and 100. The default value is 75. • (Optional) Check the Control prefixes received from the peer check box to specify additional controls for the prefixes received from a peer. Do one of the following: <ul style="list-style-type: none"> • Select Terminate peering when prefix limit is exceeded to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the Restart interval field. • Select Give only warning message when prefix limit is exceeded to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.
Routes	
Advertisement Interval	Enter the minimum interval (in seconds) between the sending of BGP routing updates. Valid values are between 1 and 600.
Remove private AS numbers from outbound routing updates	(Optional) Excludes the private AS numbers from being advertised on outbound routes.
Generate Default route	<p>(Optional) Select to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route. Enter or Select the route map that allows the route 0.0.0.0 to be injected conditionally in the Route map field.</p> <p>Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.</p>

Element	Description
Conditionally Advertised Routes	<p>(Optional) To add or edit conditionally advertised routes, click the Add Row (+) button, or select a row in the table and click the Edit Row(pencil) button.</p> <p>In the Add/Edit Advertised Route dialog box, do the following:</p> <ul style="list-style-type: none"> • Click Select to open the Route Map Object Selector from which you can select a route map that will be advertised if the conditions of the exist map or the non-exist map are met. For more information about route maps, see Understanding Route Map Objects , on page 145. • Do one of the following: <ul style="list-style-type: none"> • Select Set Exist Map and choose a route map from the Route Map Object Selector. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised. • Select Non-Exist Map and choose a route map from the Route Map Object Selector. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
Timers	
Set timers for the BGP peer	(Optional) Select to set the keepalive frequency, hold time and minimum hold time.
Keepalive Interval	Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor. Valid values are between 0 and 65535. The default value is 60 seconds.
Hold Time	Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Valid values are between 0 and 65535. The default value is 180 seconds.
Min Hold Time	(Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Valid values are between 0 and 65535. The default value is 0 seconds.
Advanced	

Element	Description
Enable Authentication	<p>(Optional) Select to enable MD5 authentication on a TCP connection between two BGP peers.</p> <ul style="list-style-type: none"> • Choose an encryption type from the Enable Encryption drop-down list. • Enter a password in the Password field. Reenter the password in the Confirm field. <p>The password is case-sensitive and can be up to 25 characters long when the service password-encryption command is enabled and up to 81 characters long when the service password-encryption command is not enabled. The first character cannot be a number. The string can contain any alphanumeric characters, including spaces.</p> <p>Note You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.</p>
Send Community attribute to this neighbor	(Optional) Specifies that communities attributes should be sent to the BGP neighbor.
Use ASA as next hop for neighbor	(Optional) Select to configure the router as the next-hop for a BGP speaking neighbor or peer group.
Disable connection verification	<p>(Optional) Select to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IP address.</p> <p>This command is required only when the neighbor ebgp-multihop command is configured with a TTL value of 1. The address of the single-hop eBGP peer must be reachable. The neighbor update-source command must be configured to allow the BGP routing process to use the loopback interface for the peering session.</p> <p>When deselected (default), a BGP routing process will verify the connection of single-hop eBGP peering session (TTL=254) to determine if the eBGP peer is directly connected to the same network segment by default. If the peer is not directly connected to same network segment, connection verification will prevent the peering session from being established.</p>
Allow connections with neighbor that is not directly connected	<p>Select to accept and attempt BGP connections to external peers residing on networks that are not directly connected.</p> <p>(Optional) Enter the time-to-live in the TTL hops field. Valid values are between 1 and 255.</p> <p>Note This feature should be used only under the guidance of Cisco technical support staff. To prevent the creation of loops through oscillating routes, the multihop will not be established if the only route to the multihop peer is the default route (0.0.0.0).</p>

Element	Description
Limit number of TTL hops to neighbor	<p>Select this option to secure a BGP peering session. Enter the maximum number of hops that separate eBGP peers in the TTL hops field. Valid values are between 1 and 254.</p> <p>This feature provides a lightweight security mechanism to protect BGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force Denial of Service (DoS) attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses in the packet headers.</p> <p>This feature leverages designed behavior of IP packets by accepting only IP packets with a TTL count that is equal to or greater than the locally configured value. Accurately forging the TTL count in an IP packet is generally considered to be impossible. Accurately forging a packet to match the TTL count from a trusted peer is not possible without internal access to the source or destination network.</p> <p>This feature should be configured on each participating router. It secures the BGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router. When this feature is enabled, BGP will establish or maintain a session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the peering session. This feature has no effect on the BGP peering session, and the peering session can still expire if keepalive packets are not received. If the TTL value in a received packet is less than the locally configured value, the packet is silently discarded and no Internet Control Message Protocol (ICMP) message is generated. This is designed behavior; a response to a forged packet is not necessary.</p> <p>To maximize the effectiveness of this feature, the hop-count value should be strictly configured to match the number of hops between the local and external network. However, you should also take path variation into account when configuring this feature for a multihop peering session.</p> <p>The following restrictions apply to the configuration of this command:</p> <ul style="list-style-type: none"> • This feature is not supported for internal BGP (iBGP) peers. • The effectiveness of this feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected peering sessions to handle the attack. • This feature is not effective against attacks from a peer that has been compromised inside of your network. This restriction also includes peers that are on the network segment between the source and destination network.
Use TCP Path MTU Discovery	(Optional) Select to enable a TCP transport session for a BGP session.
TCP transport mode	Choose the TCP connection mode from the drop-down list. Options are Default, Active, or Passive.

Element	Description
Weight	(Optional) Enter a weight for the BGP neighbor connection.
BGP Version	Choose the BGP version that the ASA will accept from the drop-down list. The version can be set to 4-Only to force the software to use only Version 4 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.
Migration	
Note	This customization should only be used for AS migration, and should be removed after the transition has been completed. The procedure should be attempted only by an experienced network operator. Routing loops can be created through improper configuration.
Customize the AS number for routes received from the neighbor	(Optional) Select to customize the AS_PATH attribute for routes received from an eBGP neighbor.
Local AS Number	Enter the local autonomous system number. Valid values are any valid autonomous system number from 1 to 4294967295 or 1.0 to 65535.65535.
Do not prepend local AS number to routes received from neighbor	(Optional) Select to prevent the local AS number from being prepended to any routes received from eBGP peer.
Replace real AS number with local AS number in routes received from neighbor	(Optional) Select to replace the real autonomous system number with the local autonomous system number in the eBGP updates. The autonomous system number from the local BGP routing process is not prepended.
Accept either real AS number or local AS number in routes received from neighbor	(Optional) Configures the eBGP neighbor to establish a peering session using the real autonomous system number (from the local BGP routing process) or by using the local autonomous system number.

Add/Edit Network Dialog Box

Use the Add/Edit Network dialog box to define the networks to be advertised by the BGP routing process.

Navigation Path

You can access the Add/Edit Network dialog box from the [IPv4 Family Tab](#) , on page 7.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv4 Family - General Tab](#) , on page 8
- [Add/Edit Aggregate Address Dialog Box](#) , on page 10
- [Add/Edit Filter Dialog Box](#) , on page 11

- [Add/Edit Neighbor Dialog Box](#) , on page 12
- [Add/Edit Redistribution Dialog Box](#) , on page 19
- [Add/Edit Route Injection Dialog Box](#) , on page 20

Field Reference

Table 7: Add/Edit Network Dialog Box

Element	Description
Network	Specifies the network to be advertised by the BGP routing processes.
Route Map	(Optional) Enter or Select a route map that should be examined to filter the networks to be advertised. If not specified, all networks are redistributed. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.

Add/Edit Redistribution Dialog Box

Use the Add/Edit Redistribution dialog box to define the conditions for redistributing routes from another routing domain into BGP.

Navigation Path

You can access the Add/Edit Redistribution dialog box from the [IPv4 Family Tab](#) , on page 7.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv4 Family - General Tab](#) , on page 8
- [Add/Edit Aggregate Address Dialog Box](#) , on page 10
- [Add/Edit Aggregate Address Dialog Box](#) , on page 24
- [Add/Edit Neighbor Dialog Box](#) , on page 12
- [Add/Edit Network Dialog Box](#) , on page 18
- [Add/Edit Route Injection Dialog Box](#) , on page 20

Field Reference

Table 8: Add/Edit Redistribution Dialog Box

Element	Description
Source Protocol	Choose the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.
Process ID	Enter the identifier for the routing process. Applies to EIGRP and OSPF routing protocols.
Metric	(Optional) Enter a metric for the redistributed route.
Route Map	<p>Enter or Select a route map that should be examined to filter the networks to be redistributed. If not specified, all networks are redistributed.</p> <p>Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.</p>
Match	<p>The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions. These options are enabled only when OSPF is chosen as the Source Protocol.</p> <ul style="list-style-type: none"> • Internal • External 1 • External 2 • NSSA External 1 • NSSA External 2

Add/Edit Route Injection Dialog Box

Use the Add/Edit Route Injection dialog box to define the routes to be conditionally injected into the BGP routing table.

Navigation Path

You can access the Add/Edit Route Injection dialog box from the [IPv4 Family Tab](#) , on page 7.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv4 Family - General Tab](#) , on page 8
- [Add/Edit Aggregate Address Dialog Box](#) , on page 10
- [Add/Edit Filter Dialog Box](#) , on page 11

- [Add/Edit Neighbor Dialog Box](#) , on page 12
- [Add/Edit Network Dialog Box](#) , on page 18
- [Add/Edit Redistribution Dialog Box](#) , on page 19

Field Reference

Table 9: Add/Edit Route Injection Dialog Box

Element	Description
Inject Map	Enter or Select the route map that specifies the prefixes to inject into the local BGP routing table. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Exist Map	Enter or Select the route map containing the prefixes that the BGP speaker will track. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Injected routes will inherit the attributes of the aggregate route	Configures the injected route to inherit attributes of the aggregate route.

IPv6 Family Tab

Use the IPv6 Family tab on the BGP page to enable and configure IPv6 settings for BGP.

Navigation Path

You can access the IPv6 Family tab from the BGP page. For more information about the BGP page, see [Configuring BGP](#) , on page 2.

Related Topics

- [About BGP](#) , on page 3
- [General Tab](#) , on page 5

Field Reference

Table 10: IPv6 Family - Aggregate Address Tab

Element	Description
Enable IPv6 Family	Enables configuration of routing sessions that use standard IPv6 address prefixes.
General	Use this panel to configure general IPv6 settings. See IPv6 Family - General Tab , on page 22 for more about these definitions.
Aggregate Address	Use this panel to define the aggregation of specific routes into one route. Specify a value for the aggregate timer (in seconds) in the Aggregate Timer field. Valid values are 0 or any value between 6 and 60. The default value is 30. See Add/Edit Aggregate Address Dialog Box , on page 24 for more about these definitions.
Neighbor	Use this panel to define BGP neighbors and neighbor settings. See Add/Edit Neighbor Dialog Box , on page 25 for more about these definitions.
Networks	Use this panel to define the networks to be advertised by the BGP routing process. See Add/Edit Network Dialog Box , on page 31 for more about these definitions.
Redistribution	Use this panel to define the conditions for redistributing routes from another routing domain into BGP. See Add/Edit Redistribution Dialog Box , on page 32 for more about these definitions.
Route Injection	Use this panel to define the routes to be conditionally injected into the BGP routing table. See Add/Edit Route Injection Dialog Box , on page 33 for more about these definitions.

IPv6 Family - General Tab

Use the IPv6 Family - General tab to configure the general IPv6 settings.

Navigation Path

You can access the General tab from the IPv6 Family Tab on the BGP page. For more information about the IPv6 Family tab, see [IPv6 Family Tab](#) , on [page 21](#).

Related Topics

- [Configuring BGP](#) , on [page 2](#)
- [About BGP](#) , on [page 3](#)

Field Reference

Table 11: IPv6 Family - General Tab

Element	Description
Scanning Interval	Enter a scanning interval (in seconds) for BGP routers for next-hop validation. Valid values are from 5 to 60 seconds. The default value is 60.
Routes and Synchronization	
Generate Default Routes	(Optional) Configures a BGP routing process to distribute a default route (network 0.0.0.0).
Advertise inactive routes	(Optional) Advertises routes that are not installed in the routing information base (RIB).
Synchronize between BGP and the Interior Gateway Protocol (IGP) system	<p>Enables synchronization between BGP and your Interior Gateway Protocol (IGP) system. To enable the Cisco IOS software to advertise a network route without waiting for the IGP, deselect this option.</p> <p>Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. By default, synchronization between BGP and the IGP is turned off to allow the Cisco IOS software to advertise a network route without waiting for route validation from the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems. Use synchronization if routers in the autonomous system do not speak BGP.</p>
Redistribute iBGP into an IGP (use filtering to limit the number of prefixes that are redistributed)	(Optional) Configures iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.
Administrative Route Distances	
External	Specifies the administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255. The default value is 20.
Internal	Specifies administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255. The default value is 200.
Local	Specifies administrative distance for local BGP routes. Local routes are those networks listed with a network router configuration command, often as back doors, for the router or for the networks that is being redistributed from another process. The range of values for this argument are from 1 to 255. The default value is 200.
Forward packets over Multiple Paths	
Number of Paths	(Optional) Specify the maximum number of Border Gateway Protocol routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.

Element	Description
IBGP Number of Paths	(Optional) Specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.

Add/Edit Aggregate Address Dialog Box

Use the Add/Edit Aggregate Address dialog box to define the aggregation of specific routes into one route.

Navigation Path

You can access the Add/Edit Aggregate Address dialog box from the [IPv6 Family Tab](#), on page 21. Click the **Add Row (+)** button, or select a row in the table and click the **Edit Row(pencil)** button.

Related Topics

- [Configuring BGP](#), on page 2
- [About BGP](#), on page 3
- [IPv6 Family - General Tab](#), on page 22
- [Add/Edit Neighbor Dialog Box](#), on page 25
- [Add/Edit Network Dialog Box](#), on page 31
- [Add/Edit Redistribution Dialog Box](#), on page 32
- [Add/Edit Route Injection Dialog Box](#), on page 33

Field Reference

Table 12: Add/Edit Aggregate Address Dialog Box

Element	Description
Network	Enter an IP address, or enter or Select the desired Network/Hosts objects.
Attribute Map	(Optional) Enter or Select the route map used to set the attribute of the aggregate route. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Advertise Map	(Optional) Enter or Select the route map used to select the routes to create AS_SET origin communities. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.

Element	Description
Suppress Map	(Optional) Enter or Select the route map used to select the routes to be suppressed. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Generate AS Set Path Information	Enables generation of autonomous system set path information.
Filter all more-specific routes from updates	Filters all more-specific routes from updates.

Add/Edit Neighbor Dialog Box

Use the Add/Edit Neighbor dialog box to define BGP neighbors and neighbor settings.

Navigation Path

You can access the Add/Edit Neighbor dialog box from the [IPv6 Family Tab](#) , on page 21. Click the **Add Row (+)** button, or select a row in the table and click the **Edit Row(pencil)** button.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv6 Family - General Tab](#) , on page 22
- [Add/Edit Aggregate Address Dialog Box](#) , on page 24
- [Add/Edit Network Dialog Box](#) , on page 31
- [Add/Edit Redistribution Dialog Box](#) , on page 32
- [Add/Edit Route Injection Dialog Box](#) , on page 33

Field Reference

Table 13: Add/Edit Neighbor Dialog Box

Element	Description
General	
IP Address	Enter the BGP neighbor IP address. This IP address is added to the BGP neighbor table.
Remote AS	Enter the autonomous system to which the BGP neighbor belongs.
Enable Address Family	(Optional) Enables communication with the BGP neighbor.

Element	Description
Shutdown neighbor administratively	(Optional) Disable a neighbor or peer group.
Configure Graceful Restart per neighbor (ASA 9.3.1+ only)	(Optional) Enables configuration of the Border Gateway Protocol (BGP) graceful restart capability for this neighbor. After selecting this option, you must use the Graceful Restart (Use in failover or spanned cluster mode) option to specify whether graceful restart should be enabled or disabled for this neighbor.
Graceful Restart (Use in failover or spanned cluster mode) (ASA 9.3.1+ only)	(Optional) Enables the Border Gateway Protocol (BGP) graceful restart capability for this neighbor.
Description	(Optional) Enter a description for the BGP neighbor.
fall-over BFD	(Optional) Enables BFD support for fall-over for the BGP neighbor.
BFD-Hop	(Optional) Specify if there is a single IP hop or multiple IP hops between a BFD source and destination.
Filtering	
Filter routes using an access list	(Optional) Enter or Select the appropriate incoming or outgoing access control list to distribute BGP neighbor information.
Filter routes using route map	(Optional) Enter or Select the appropriate incoming or outgoing route maps to apply a route map to incoming or outgoing routes. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Filter routes using a Prefix list	(Optional) Enter or Select the appropriate incoming or outgoing prefix list to distribute BGP neighbor information. Tip Click Select to open the Prefix List Object Selector from which you can select a prefix list object. You can also create new objects from the object Prefix List Object selector. For more information, see Add or Edit Prefix List Object Dialog Box , on page 159.
Filter routes using AS Path filter	(Optional) Enter or Select the appropriate incoming or outgoing AS path filter to distribute BGP neighbor information. Tip Click Select to open the AS Path Object Selector from which you can select an AS path object. You can also create new AS path objects from the AS Path Object Selector. For more information, see Add or Edit As Path Object Dialog Boxes , on page 163.

Element	Description
Limit the number of prefixes allowed from the neighbor	<p>(Optional) Select to control the number of prefixes that can be received from a neighbor.</p> <ul style="list-style-type: none"> • Enter the maximum number of prefixes allowed from a specific neighbor in the Maximum Prefixes field. • Enter the percentage (of maximum) at which the router starts to generate a warning message in the Threshold Level field. Valid values are integers between 1 and 100. The default value is 75. • (Optional) Check the Control prefixes received from the peer check box to specify additional controls for the prefixes received from a peer. Do one of the following: <ul style="list-style-type: none"> • Select Terminate peering when prefix limit is exceeded to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the Restart interval field. • Select Give only warning message when prefix limit is exceeded to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.
Routes	
Advertisement Interval	Enter the minimum interval (in seconds) between the sending of BGP routing updates. Valid values are between 1 and 600.
Remove private AS numbers from outbound routing updates	(Optional) Excludes the private AS numbers from being advertised on outbound routes.
Generate Default route	<p>(Optional) Select to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route. Enter or Select the route map that allows the route 0.0.0.0 to be injected conditionally in the Route map field.</p> <p>Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.</p>

Element	Description
Conditionally Advertised Routes	<p>(Optional) To add or edit conditionally advertised routes, click the Add Row (+) button, or select a row in the table and click the Edit Row(pencil) button.</p> <p>In the Add/Edit Advertised Route dialog box, do the following:</p> <ul style="list-style-type: none"> • Click Select to open the Route Map Object Selector from which you can select a route map that will be advertised if the conditions of the exist map or the non-exist map are met. For more information about route maps, see Understanding Route Map Objects , on page 145. • Do one of the following: <ul style="list-style-type: none"> • Select Set Exist Map and choose a route map from the Route Map Object Selector. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised. • Select Non-Exist Map and choose a route map from the Route Map Object Selector. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
Timers	
Set timers for the BGP peer	(Optional) Select to set the keepalive frequency, hold time and minimum hold time.
Keepalive Interval	Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor. Valid values are between 0 and 65535. The default value is 60 seconds.
Hold Time	Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Valid values are between 0 and 65535. The default value is 180 seconds.
Min Hold Time	(Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Valid values are between 0 and 65535. The default value is 0 seconds.
Advanced	

Element	Description
Enable Authentication	<p>(Optional) Select to enable MD5 authentication on a TCP connection between two BGP peers.</p> <ul style="list-style-type: none"> • Choose an encryption type from the Enable Encryption drop-down list. • Enter a password in the Password field. Reenter the password in the Confirm field. <p>The password is case-sensitive and can be up to 25 characters long when the service password-encryption command is enabled and up to 81 characters long when the service password-encryption command is not enabled. The first character cannot be a number. The string can contain any alphanumeric characters, including spaces.</p> <p>Note You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.</p>
Send Community attribute to this neighbor	<p>(Optional) Specifies that communities attributes should be sent to the BGP neighbor.</p>
Use ASA as next hop for neighbor	<p>(Optional) Select to configure the router as the next-hop for a BGP speaking neighbor or peer group.</p>
Disable connection verification	<p>(Optional) Select to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IP address.</p> <p>This command is required only when the neighbor ebgp-multihop command is configured with a TTL value of 1. The address of the single-hop eBGP peer must be reachable. The neighbor update-source command must be configured to allow the BGP routing process to use the loopback interface for the peering session.</p> <p>When deselected (default), a BGP routing process will verify the connection of single-hop eBGP peering session (TTL=254) to determine if the eBGP peer is directly connected to the same network segment by default. If the peer is not directly connected to same network segment, connection verification will prevent the peering session from being established.</p>
Allow connections with neighbor that is not directly connected	<p>Select to accept and attempt BGP connections to external peers residing on networks that are not directly connected.</p> <p>(Optional) Enter the time-to-live in the TTL hops field. Valid values are between 1 and 255.</p> <p>Note This feature should be used only under the guidance of Cisco technical support staff. To prevent the creation of loops through oscillating routes, the multihop will not be established if the only route to the multihop peer is the default route (0.0.0.0).</p>

Element	Description
Limit number of TTL hops to neighbor	<p>Select this option to secure a BGP peering session. Enter the maximum number of hops that separate eBGP peers in the TTL hops field. Valid values are between 1 and 254.</p> <p>This feature provides a lightweight security mechanism to protect BGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force Denial of Service (DoS) attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses in the packet headers.</p> <p>This feature leverages designed behavior of IP packets by accepting only IP packets with a TTL count that is equal to or greater than the locally configured value. Accurately forging the TTL count in an IP packet is generally considered to be impossible. Accurately forging a packet to match the TTL count from a trusted peer is not possible without internal access to the source or destination network.</p> <p>This feature should be configured on each participating router. It secures the BGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router. When this feature is enabled, BGP will establish or maintain a session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the peering session. This feature has no effect on the BGP peering session, and the peering session can still expire if keepalive packets are not received. If the TTL value in a received packet is less than the locally configured value, the packet is silently discarded and no Internet Control Message Protocol (ICMP) message is generated. This is designed behavior; a response to a forged packet is not necessary.</p> <p>To maximize the effectiveness of this feature, the hop-count value should be strictly configured to match the number of hops between the local and external network. However, you should also take path variation into account when configuring this feature for a multihop peering session.</p> <p>The following restrictions apply to the configuration of this command:</p> <ul style="list-style-type: none"> • This feature is not supported for internal BGP (iBGP) peers. • The effectiveness of this feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected peering sessions to handle the attack. • This feature is not effective against attacks from a peer that has been compromised inside of your network. This restriction also includes peers that are on the network segment between the source and destination network.
Use TCP Path MTU Discovery	(Optional) Select to enable a TCP transport session for a BGP session.
TCP transport mode	Choose the TCP connection mode from the drop-down list. Options are Default, Active, or Passive.

Element	Description
Weight	(Optional) Enter a weight for the BGP neighbor connection.
BGP Version	Choose the BGP version that the ASA will accept from the drop-down list. The version can be set to 4-Only to force the software to use only Version 4 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.
Migration	
Note	This customization should only be used for AS migration, and should be removed after the transition has been completed. The procedure should be attempted only by an experienced network operator. Routing loops can be created through improper configuration.
Customize the AS number for routes received from the neighbor	(Optional) Select to customize the AS_PATH attribute for routes received from an eBGP neighbor.
Local AS Number	Enter the local autonomous system number. Valid values are any valid autonomous system number from 1 to 4294967295 or 1.0 to 65535.65535.
Do not prepend local AS number to routes received from neighbor	(Optional) Select to prevent the local AS number from being prepended to any routes received from eBGP peer.
Replace real AS number with local AS number in routes received from neighbor	(Optional) Select to replace the real autonomous system number with the local autonomous system number in the eBGP updates. The autonomous system number from the local BGP routing process is not prepended.
Accept either real AS number or local AS number in routes received from neighbor	(Optional) Configures the eBGP neighbor to establish a peering session using the real autonomous system number (from the local BGP routing process) or by using the local autonomous system number.

Add/Edit Network Dialog Box

Use the Add/Edit Network dialog box to define the networks to be advertised by the BGP routing process.

Navigation Path

You can access the Add/Edit Network dialog box from the [IPv6 Family Tab](#), on page 21. Click the **Add Row (+)** button, or select a row in the table and click the **Edit Row(pencil)** button.

Related Topics

- [Configuring BGP](#), on page 2
- [About BGP](#), on page 3
- [IPv6 Family - General Tab](#), on page 22
- [Add/Edit Aggregate Address Dialog Box](#), on page 24

- [Add/Edit Neighbor Dialog Box](#) , on page 25
- [Add/Edit Redistribution Dialog Box](#) , on page 32
- [Add/Edit Route Injection Dialog Box](#) , on page 33

Field Reference

Table 14: Add/Edit Network Dialog Box

Element	Description
Network	Specifies the network to be advertised by the BGP routing processes.
Route Map	(Optional) Enter or Select a route map that should be examined to filter the networks to be advertised. If not specified, all networks are redistributed. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.

Add/Edit Redistribution Dialog Box

Use the Add/Edit Redistribution dialog box to define the conditions for redistributing routes from another routing domain into BGP.

Navigation Path

You can access the Add/Edit Redistribution dialog box from the [IPv6 Family Tab](#) , on page 21. Click the **Add Row (+)** button, or select a row in the table and click the **Edit Row(pencil)** button.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv6 Family - General Tab](#) , on page 22
- [Add/Edit Aggregate Address Dialog Box](#) , on page 24
- [Add/Edit Neighbor Dialog Box](#) , on page 25
- [Add/Edit Network Dialog Box](#) , on page 31
- [Add/Edit Route Injection Dialog Box](#) , on page 33

Field Reference

Table 15: Add/Edit Redistribution Dialog Box

Element	Description
Source Protocol	Choose the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.

Element	Description
Process ID	Enter the identifier for the routing process. Applies to EIGRP and OSPF routing protocols.
Metric	(Optional) Enter a metric for the redistributed route.
Route Map	<p>Enter or Select a route map that should be examined to filter the networks to be redistributed. If not specified, all networks are redistributed.</p> <p>Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.</p>
Match	<p>The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions. These options are enabled only when OSPF is chosen as the Source Protocol.</p> <ul style="list-style-type: none"> • Internal • External 1 • External 2 • NSSA External 1 • NSSA External 2

Add/Edit Route Injection Dialog Box

Use the Add/Edit Route Injection dialog box to define the routes to be conditionally injected into the BGP routing table.

Navigation Path

You can access the Add/Edit Route Injection dialog box from the [IPv6 Family Tab](#) , on page 21. Click the **Add Row (+)** button, or select a row in a table and click the **Edit Row (pencil)** button.

Related Topics

- [Configuring BGP](#) , on page 2
- [About BGP](#) , on page 3
- [IPv6 Family - General Tab](#) , on page 22
- [Add/Edit Aggregate Address Dialog Box](#) , on page 24
- [Add/Edit Neighbor Dialog Box](#) , on page 25
- [Add/Edit Network Dialog Box](#) , on page 31
- [Add/Edit Redistribution Dialog Box](#) , on page 32

Field Reference

Table 16: Add/Edit Route Injection Dialog Box

Element	Description
Inject Map	Enter or Select the route map that specifies the prefixes to inject into the local BGP routing table. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Exist Map	Enter or Select the route map containing the prefixes that the BGP speaker will track. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Injected routes will inherit the attributes of the aggregate route	Configures the injected route to inherit attributes of the aggregate route.

Configuring EIGRP

The EIGRP page provides six tabbed panels for configuring Enhanced Interior Gateway Routing Protocol (EIGRP) routing on a firewall device. The following topics provide detailed information about enabling and configuring EIGRP:

- [About EIGRP](#) , on page 35
- [EIGRP Advanced Dialog Box](#) , on page 36
- [Setup Tab](#) , on page 38
- [Filter Rules Tab](#) , on page 41
- [Neighbors Tab](#) , on page 42
- [Redistribution Tab](#) , on page 44
- [Summary Address Tab](#) , on page 47
- [Interfaces Tab](#) , on page 49

Navigation Path

- (Device view) Select **Platform** > **Routing** > **EIGRP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform** > **Routing** > **EIGRP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 17: EIGRP Page

Element	Description
Enable EIGRP	Check this box to enable the EIGRP routing process.
AS Number	Enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.
Advanced button	Opens the EIGRP Advanced Dialog Box , on page 36, in which you can configure additional EIGRP process settings, such as the router ID, stub routing, and adjacency changes.
Setup tab	Use the Setup tab to configure the networks used by the EIGRP routing process, passive interfaces, default route information, administrative distances, and default metrics. For more information, see Setup Tab , on page 38.
Filter Rules tab	Use the Filter Rules tab to define filter rules that let you control which routes are accepted or advertised by the EIGRP routing process. For more information, see Filter Rules Tab , on page 41.
Neighbors tab	Use the Neighbors tab to manually define EIGRP neighbors. For more information, see Neighbors Tab , on page 42.
Redistribution tab	Use the Redistribution tab to define the rules for redistributing routes from other routing protocols to the EIGRP routing process. For more information, see Redistribution Tab , on page 44.
Summary Address tab	Use the Summary Address tab to create statically defined EIGRP summary addresses. For more information, see Summary Address Tab , on page 47.
Interfaces tab	Use the Interfaces tab to configure interfaces for EIGRP. For more information, see Interfaces Tab , on page 49.

About EIGRP

EIGRP is an enhanced version of IGRP developed by Cisco. Unlike IGRP and RIP, EIGRP does not send out periodic route updates. EIGRP updates are sent out only when the network topology changes. Key capabilities that distinguish EIGRP from other routing protocols include fast convergence, support for variable-length subnet mask, support for partial updates, and support for multiple network layer protocols.

A router running EIGRP stores all the neighbor routing tables so that it can quickly adapt to alternate routes. If no appropriate route exists, EIGRP queries its neighbors to discover an alternate route. These queries propagate until an alternate route is found. Its support for variable-length subnet masks permits routes to be automatically summarized on a network number boundary. In addition, EIGRP can be configured to summarize on any bit boundary at any interface. EIGRP does not make periodic updates. Instead, it sends partial updates only when the metric for a route changes. Propagation of partial updates is automatically bounded so that only

those routers that need the information are updated. As a result of these two capabilities, EIGRP consumes significantly less bandwidth than IGRP.

Neighbor discovery is the process that the ASA uses to dynamically learn of other routers on directly attached networks. EIGRP routers send out multicast hello packets to announce their presence on the network. When the ASA receives a hello packet from a new neighbor, it sends its topology table to the neighbor with an initialization bit set. When the neighbor receives the topology update with the initialization bit set, the neighbor sends its topology table back to the ASA.

The hello packets are sent out as multicast messages. No response is expected to a hello message. The exception to this is for statically defined neighbors. If you manually configure a neighbor, the hello messages sent to that neighbor are sent as unicast messages. Routing updates and acknowledgments are sent out as unicast messages.

Once this neighbor relationship is established, routing updates are not exchanged unless there is a change in the network topology. The neighbor relationship is maintained through the hello packets. Each hello packet received from a neighbor includes a hold time. This is the time in which the ASA can expect to receive a hello packet from that neighbor. If the ASA does not receive a hello packet from that neighbor within the hold time advertised by that neighbor, the ASA considers that neighbor to be unavailable.

The EIGRP protocol uses four key algorithm technologies, four key technologies, including neighbor discovery/recovery, Reliable Transport Protocol (RTP), and DUAL, which is important for route computations. DUAL saves all routes to a destination in the topology table, not just the least-cost route. The least-cost route is inserted into the routing table. The other routes remain in the topology table. If the main route fails, another route is chosen from the feasible successors. A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination. The feasibility calculation guarantees that the path is not part of a routing loop.

If a feasible successor is not found in the topology table, a route recomputation must occur. During route recomputation, DUAL queries the EIGRP neighbors for a route, who in turn query their neighbors. Routers that do not have a feasible successor for the route return an unreachable message.

During route recomputation, DUAL marks the route as active. By default, the ASA waits for three minutes to receive a response from its neighbors. If the ASA does not receive a response from a neighbor, the route is marked as stuck-in-active. All routes in the topology table that point to the unresponsive neighbor as a feasibility successor are removed.



Note EIGRP neighbor relationships are not supported through the IPsec tunnel without a GRE tunnel.

Related Topics

- [Configuring EIGRP](#) , on page 34

EIGRP Advanced Dialog Box

Use the EIGRP Advanced dialog box to configure settings such as the router ID, stub routing, and adjacency changes.

Navigation Path

You can access the EIGRP Advanced dialog box from the EIGRP page (see [Configuring EIGRP](#), on page 34).

Related Topics

- [Configuring EIGRP](#), on page 34

Field Reference

Table 18: EIGRP Advanced Dialog Box

Element	Description
Router ID	<p>The router ID is used to identify the originating router for external routes. If an external route is received with the local router ID, the route is discarded. To prevent this, specify a global address for the router ID. A unique value should be configured for each EIGRP router.</p> <p>On a single device, choose Automatic or IP Address. (An address field appears when you choose IP Address.)</p> <p>If you choose Automatic, the highest-level IP address on the security appliance is used as the router ID. To use a fixed router ID, choose IP Address and enter an IPv4 address in the Router ID field.</p> <p>On a device cluster, choose Automatic or Cluster Pool. (An IPv4 Pool object ID field appears when you choose Cluster Pool.)</p> <p>If you choose Cluster Pool, enter or Select the name of the IPv4 Pool object that is to supply the Router ID address. For more information, see .</p>

Element	Description
Stub	<p>You can enable, and configure the ASA as an EIGRP stub router. Stub routing decreases memory and processing Add or Edit IPv4 Pool Dialog Box requirements on the ASA. As a stub router, the ASA does not need to maintain a complete EIGRP routing table because it forwards all nonlocal traffic to a distribution router. Generally, the distribution router need not send anything more than a default route to the stub router.</p> <p>Only specified routes are propagated from the stub router to the distribution router. As a stub router, the ASA responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message “inaccessible.” When the ASA is configured as a stub, it sends a special peer information packet to all neighboring routers to report its status as a stub router. Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router depends on the distribution router to send the correct updates to all peers.</p> <p>To enable the ASA as an EIGRP stub routing process, choose one or more of the following EIGRP stub routing processes:</p> <ul style="list-style-type: none"> • Receive only—Configures the EIGRP stub routing process to receive route information from the neighbor routers but does not send route information to the neighbors. If this option is selected, you cannot select any of the other stub routing options. • Connected—Advertises connected routes. • Redistributed—Advertises redistributed routes. • Static—Advertises static routes. • Summary—Advertises summary routes.
Adjacency Changes	<p>These options specify the syslog messages sent when adjacency changes occur.</p> <ul style="list-style-type: none"> • Log Neighbor Changes – enables the logging of EIGRP neighbor adjacency changes. This option is selected by default. • Log Neighbor Warnings – enables the logging of EIGRP neighbor warning messages. This option is selected by default. <p>(Optional) The time interval (in seconds) between repeated neighbor warning messages. Valid values are from 1 to 65535. Repeated warnings are not logged if they occur during this interval.</p>

Setup Tab

Use the Setup tab on the EIGRP page to configure the networks used by the EIGRP routing process, passive interfaces, default route information, administrative distances, and default metrics.

Navigation Path

You can access the Setup tab from the EIGRP Page; see [Configuring EIGRP](#), on page 34 for more information.

Related Topics

- [Configuring EIGRP](#) , on page 34
- [About EIGRP](#) , on page 35
- [Setup Tab](#) , on page 38
- [Filter Rules Tab](#) , on page 41
- [Neighbors Tab](#) , on page 42
- [Redistribution Tab](#) , on page 44
- [Summary Address Tab](#) , on page 47
- [Interfaces Tab](#) , on page 49

Field Reference

Table 19: EIGRP - Setup Tab

Element	Description
Auto Summary	<p>Check this box to enable automatic route summarization. Auto summary is enabled by default for ASA versions earlier than 9.2.1 and is disabled by default for ASA 9.2(1) and later.</p> <p>When enabled, the EIGRP routing process summarizes on network number boundaries. This can cause routing problems if you have noncontiguous networks.</p> <p>For example, if you have a router with the networks 192.168.1.0, 192.168.2.0, and 192.168.3.0 connected to it, and those networks all participate in EIGRP, the EIGRP routing process creates the summary address 192.168.0.0 for those routes. If an additional router is added to the network with the networks 192.168.10.0 and 192.168.11.0, and those networks participate in EIGRP, they will also be summarized as 192.168.0.0. To prevent the possibility of traffic being routed to the wrong location, you should disable automatic route summarization on the routers creating the conflicting summary addresses.</p>
Networks	<p>Enter the IP addresses of the networks to participate in the EIGRP routing process.</p> <p>Tip You can click Select to select the networks from a list of network/host objects.</p>
Passive Interface	<p>You can configure one or more interfaces as passive interfaces. In EIGRP, a passive interface does not send or receive routing updates.</p> <p>By default, all interfaces are enabled for active routing (sending and receiving routing updates) when routing is enabled for that interface.</p> <p>To configure passive interfaces, do one of the following:</p> <ul style="list-style-type: none"> • To enable all interfaces for active routing (sending and receiving routing updates) when routing is enabled for that interface, select None. • To configure all interfaces as passive, select All Interfaces. • To configure specific interfaces as passive, select Specified Interfaces and then enter or select the interfaces that you want to make passive.

Element	Description
Default Route Information	<p>You can control the sending and receiving of default route information in EIGRP updates. By default, default routes are sent and accepted. Configuring the ASA to disallow default information to be received causes the candidate default route bit to be blocked on received routes. Configuring the ASA to disallow default information to be sent disables the setting of the default route bit in advertised routes.</p> <ul style="list-style-type: none"> • Accept Default Route Info—configures EIGRP to accept exterior default routing information. Optionally, you can specify a standard access list that define which networks are allowed and which are not when receiving default route information. • Send Default Route Info—configures EIGRP to advertise external routing information. Optionally, you can specify a standard access list that defines which networks are allowed and which are not when sending default route information.
Administrative Distance	<p>Because every routing protocol has metrics based on algorithms that are different from the other routing protocols, it is not always possible to determine the “best path” for two routes to the same destination that were generated by different routing protocols. Administrative distance is a route parameter that the ASA uses to select the best path when there are two or more different routes to the same destination from two different routing protocols.</p> <p>If you have more than one routing protocol running on the ASA, you can use the distance eigrp command to adjust the default administrative distances of routes discovered by the EIGRP routing protocol in relation to the other routing protocols.:</p> <p>Internal Distance—Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system. Valid values are from 1 to 255. The default value is 90.</p> <p>External Distance—Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a neighbor external to the autonomous system. Valid values are from 1 to 255. The default value is 170.</p>
Default Metrics	<p>You can define the default metrics for routes redistributed into the EIGRP routing process:</p> <ul style="list-style-type: none"> • Bandwidth—the minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295. • Delay Time—the route delay in tens of microseconds. Valid values range from 0 to 4294967295. • Reliability—the likelihood of successful packet transmission expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability. • Loading—the effective bandwidth of the route. Valid values range from 1 to 255; 255 indicates 100 percent loaded. • MTU—the smallest allowed value for the maximum transmission unit of the path. Valid values range from 1 to 65535.

Filter Rules Tab

The Filter Rules tab contains the Filter Rules table which displays the route filtering rules configured for the EIGRP routing process. Filter rules let you control which routes are accepted or advertised by the EIGRP routing process.

Navigation Path

You can access the Filter Rules tab from the EIGRP Page; see [Configuring EIGRP , on page 34](#) for more information.

Related Topics

- [Add/Edit EIGRP Filter Rule Dialog Box , on page 41](#)
- [Configuring EIGRP , on page 34](#)
- [About EIGRP , on page 35](#)
- [Setup Tab , on page 38](#)
- [Neighbors Tab , on page 42](#)
- [Redistribution Tab , on page 44](#)
- [Summary Address Tab , on page 47](#)
- [Interfaces Tab , on page 49](#)

Field Reference

Table 20: EIGRP - Filter Rules Tab

Element	Description
Direction	The direction for the filter rule: <ul style="list-style-type: none"> • Inbound—The rule filters default route information from incoming EIGRP routing updates. • Outbound—The rule filters default route information from outgoing EIGRP routing updates.
Interface	(Optional) The interface to which the filter rule applies.
Protocol	The routing protocol being filtered: BGP, Connected, OSPF, RIP, or Static.
ACL	Standard IP access list name. The list defines which networks are to be received and which are to be suppressed in routing updates.

Add/Edit EIGRP Filter Rule Dialog Box

Use the Add/Edit EIGRP Filter Rule dialog box to add new filter rules to the Filter Rules table or to modify an existing filter rule.

Navigation Path

You can access the Add/Edit EIGRP Filter Rule dialog box from the [Filter Rules Tab](#), on page 41.

Related Topics

- [Configuring EIGRP](#), on page 34
- [About EIGRP](#), on page 35
- [Filter Rules Tab](#), on page 41

Field Reference

Table 21: Add/Edit EIGRP Filter Rule Dialog Box

Element	Description
EIGRP Filter Direction	Specify the direction for the filter rule: <ul style="list-style-type: none"> • Inbound—The rule filters default route information from incoming EIGRP routing updates. • Outbound—The rule filters default route information from outgoing EIGRP routing updates.
Type	Specify the type of filter rule: <ul style="list-style-type: none"> • (Optional) Interface—Specify the interface on which to apply the routing updates. Specifying an interface causes the access list to be applied only to routing updates for that interface. If no interface is specified, the access list will be applied to all updates. • (Optional) Routing Protocol—For outbound EIGRP routing updates, select the routing protocol for which you want to filter: BGP, Connected, OSPF, RIP, or Static. Routing Protocol ID—Enter the identifier for the routing process. Applies to BGP and OSPF routing protocols.
ACL	Select an Access Control List that defines which networks are to be received and which are to be suppressed in routing updates.

Neighbors Tab

The Neighbors tab contains the Neighbors table, through which you can define static neighbors. When you manually define an EIGRP neighbor, hello packets are sent to that neighbor as unicast messages.

Navigation Path

You can access the Neighbors tab from the EIGRP Page; see [Configuring EIGRP](#), on page 34 for more information.

Related Topics

- [Add/Edit EIGRP Neighbor Dialog Box](#) , on page 43
- [Configuring EIGRP](#) , on page 34
- [About EIGRP](#) , on page 35
- [Setup Tab](#) , on page 38
- [Filter Rules Tab](#) , on page 41
- [Redistribution Tab](#) , on page 44
- [Summary Address Tab](#) , on page 47
- [Interfaces Tab](#) , on page 49

Field Reference

Table 22: EIGRP - Neighbors Tab

Element	Description
Interface	The interface through which the neighbor is available.
Neighbor	The IP address of the static neighbor.

Add/Edit EIGRP Neighbor Dialog Box

EIGRP hello packets are sent as multicast packets. If an EIGRP neighbor is located across a non broadcast network, such as a tunnel, you must manually define that neighbor. When you manually define an EIGRP neighbor, hello packets are sent to that neighbor as unicast messages.



Note Configuring the passive-interface command for an interface suppresses all incoming and outgoing routing updates and hello messages on that interface. EIGRP neighbor adjacencies cannot be established or maintained over an interface that is configured as passive.

Use the Add/Edit EIGRP Neighbor dialog box to define a static neighbor or change information for an existing static neighbor.

Navigation Path

You can access the Add/Edit EIGRP Neighbor dialog box from the [Neighbors Tab](#) , on page 42.

Related Topics

- [Configuring EIGRP](#) , on page 34
- [About EIGRP](#) , on page 35
- [Neighbors Tab](#) , on page 42

Field Reference

Table 23: Add/Edit EIGRP Neighbor Dialog Box

Element	Description
Interface	The interface through which the neighbor is available. Tip You can click Select to select the interface from a list of interface objects.
Neighbor	The IP address of the static neighbor. Tip You can click Select to select the neighbor from a list of host objects.

Redistribution Tab

Use the Redistribution tab to define the rules for redistributing routes from other routing protocols to the EIGRP routing process.

Navigation Path

You can access the Redistribution tab from the EIGRP Page; see [Configuring EIGRP](#), on page 34 for more information.

Related Topics

- [Add/Edit EIGRP Redistribution Dialog Box](#), on page 45
- [Configuring EIGRP](#), on page 34
- [About EIGRP](#), on page 35
- [Setup Tab](#), on page 38
- [Filter Rules Tab](#), on page 41
- [Neighbors Tab](#), on page 42
- [Summary Address Tab](#), on page 47
- [Interfaces Tab](#), on page 49

Field Reference

Table 24: EIGRP - Redistribution Tab

Element	Description
Protocol	<p>The source protocol from which the routes are being redistributed:</p> <ul style="list-style-type: none"> • BGP—Redistribute routes discovered by the BGP routing process to EIGRP. • RIP—Redistributes routes discovered by the RIP routing process to EIGRP. • Static—Redistributes static routes to the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them. • Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them. • OSPF—Redistributes routes discovered by the OSPF routing process to EIGRP. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, or BGP routes.
ID	The autonomous system (AS) number for the BGP or OSPF routing process.
Bandwidth	The minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295.
Delay Time	The route delay in tens of microseconds. Valid values range from 0 to 4294967295.
Reliability	The likelihood of successful packet transmission expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability.
Loading	The effective bandwidth of the route. Valid values range from 1 to 255; 255 indicates 100 percent loaded.
MTU	The smallest allowed value for the maximum transmission unit of the path. Valid values range from 1 to 65535.
Route Map	The name of the route map object to apply to the redistribution entry.

Add/Edit EIGRP Redistribution Dialog Box

Use the Add/Edit Redistribution dialog box to add a redistribution rule or to edit an existing redistribution rule in the Redistribution table.

Navigation Path

You can access the Add/Edit EIGRP Redistribution dialog box from the [Redistribution Tab](#) , on page 44.

Related Topics

- [Configuring EIGRP](#) , on page 34
- [About EIGRP](#) , on page 35
- [Redistribution Tab](#) , on page 44

Field Reference**Table 25: Add/Edit EIGRP Redistribution Dialog Box**

Element	Description
Protocol	<p>Select the source protocol from which the routes are being redistributed. You can choose one of the following options:</p> <ul style="list-style-type: none"> • BGP—Redistribute routes discovered by the BGP routing process to EIGRP. • RIP—Redistributes routes discovered by the RIP routing process to EIGRP. • Static—Redistributes static routes to the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them. • Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them. • OSPF—Redistributes routes discovered by the OSPF routing process to EIGRP. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, or BGP routes.
Routing Process ID	The autonomous system (AS) number for the BGP or OSPF routing process.
Optional Metrics	<p>You can define the following metrics for routes redistributed into the EIGRP routing process:</p> <ul style="list-style-type: none"> • Bandwidth—the minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295. • Delay Time—the route delay in tens of microseconds. Valid values range from 0 to 4294967295. • Reliability—the likelihood of successful packet transmission expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability. • Loading—the effective bandwidth of the route. Valid values range from 1 to 255; 255 indicates 100 percent loaded. • MTU—the smallest allowed value for the maximum transmission unit of the path. Valid values range from 1 to 65535.

Element	Description
Route Map	<p>Enter or Select a route map object to define which routes are redistributed into the EIGRP routing process.</p> <p>Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.</p>
Optional OSPF Redistribution	<p>If you have chosen OSPF as the Route Type, choose the conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions:</p> <ul style="list-style-type: none"> • Internal—The route is internal to a specific AS. • External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes. • External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes. • NSSA External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes. • NSSA External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.

Summary Address Tab

Use the Summary Address tab to configure a summary for EIGRP on a specific interface. You can configure summary addresses on a per-interface basis. You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on an ASA with automatic route summarization disabled. If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

Navigation Path

You can access the Summary Address tab from the EIGRP Page; see [Configuring EIGRP](#) , on page 34 for more information.

Related Topics

- [Add/Edit EIGRP Summary Address Dialog Box](#) , on page 48
- [Configuring EIGRP](#) , on page 34
- [About EIGRP](#) , on page 35
- [Setup Tab](#) , on page 38
- [Filter Rules Tab](#) , on page 41
- [Neighbors Tab](#) , on page 42

- [Redistribution Tab](#) , on page 44
- [Interfaces Tab](#) , on page 49

Field Reference

Table 26: EIGRP - Summary Address Tab

Element	Description
Interface	The interface from which the summary address is advertised.
Network	The IP address and network mask of the summary address.
Administrative Distance	The administrative distance of the summary route.

Add/Edit EIGRP Summary Address Dialog Box

Use the Add/Edit EIGRP Summary Address dialog box to add new entries or to modify existing entries in the Summary Address table. You can configure summary addresses on a per-interface basis. You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on an ASA with automatic route summarization disabled. If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

Navigation Path

You can access the Add/Edit EIGRP Summary Address dialog box from the [Summary Address Tab](#) , on page 47.

Related Topics

- [Configuring EIGRP](#) , on page 34
- [About EIGRP](#) , on page 35
- [Summary Address Tab](#) , on page 47

Field Reference

Table 27: Add/Edit EIGRP Summary Address Dialog Box

Element	Description
Interface	The interface from which the summary address is advertised. Tip You can click Select to select the interface from a list of interface objects.
Networks	The IP address and network mask of the summary address. Tip You can click Select to select the network from a list of network objects.
Administrative Distance	(Optional) The administrative distance of the summary route. Valid values are from 1 to 255. The default value is 5.

Interfaces Tab

Use the Interfaces tab to configure interface-specific EIGRP routing properties.

Navigation Path

You can access the Interfaces tab from the EIGRP Page; see [Configuring EIGRP](#) , on page 34 for more information.

Related Topics

- [Add/Edit EIGRP Interface Dialog Box](#) , on page 49
- [Configuring EIGRP](#) , on page 34
- [About EIGRP](#) , on page 35
- [Setup Tab](#) , on page 38
- [Filter Rules Tab](#) , on page 41
- [Neighbors Tab](#) , on page 42
- [Redistribution Tab](#) , on page 44
- [Summary Address Tab](#) , on page 47

Field Reference

Table 28: EIGRP - Interfaces Tab

Element	Description
Interface	The name of the interface to which the configuration applies.
Hello Interval	The interval, in seconds, between EIGRP hello packets sent on an interface. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Hold Time	The hold time advertised by the ASA in EIGRP hello packets. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.
Split Horizon	Whether EIGRP split-horizon is enabled (true) or disabled (false) on an interface.
Delay	The delay time in tens of microseconds. Valid values are from 1 to 16777215. This option is not supported for devices in multi-context mode.
Key ID	The ID of the key used to authenticate EIGRP updates.

Add/Edit EIGRP Interface Dialog Box

Use the Add/Edit EIGRP Interface dialog box to configure interface-specific EIGRP routing parameters.

Navigation Path

You can access the Add/Edit EIGRP Interface dialog box from the [Interfaces Tab](#) , on page 49.

Related Topics

- [Configuring EIGRP](#) , on page 34
- [About EIGRP](#) , on page 35
- [Interfaces Tab](#) , on page 49

Field Reference*Table 29: Add/Edit EIGRP Interface Dialog Box*

Element	Description
Interface	The name of the interface to which the configuration applies.
Hello Interval	The interval, in seconds, between EIGRP hello packets sent on an interface. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Hold Time	The hold time advertised by the ASA in EIGRP hello packets. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.
Split Horizon	Enable/disable EIGRP split-horizon on an interface.
Delay Time	The delay time in tens of microseconds. Valid values are from 1 to 16777215. This option is not supported for devices in multi-context mode and will be disabled.
Enable MD5 Authentication	Enables MD5 authentication of EIGRP packets.
Key Type	Select Clear Text to indicate that the key you will be entering is in clear text. Select Encrypted to indicate that the key you will be entering is already encrypted.
Key ID and Key	Specify the key to authenticate EIGRP updates: <ul style="list-style-type: none"> • Key ID—Enter a numerical key identifier. Valid values range from 0 to 255. • Key—An alphanumeric character string of up to 16 bytes. • Confirm—Re-enter the key.

Configuring ISIS

The ISIS page provides nine tabbed panels for configuring ISIS (Intermediate System-to-Intermediate System) routing on a firewall device. ISIS routing protocol is supported from Security Manager version 4.11 for ASA devices running the software version 9.6(1) or later. The following topics provide detailed information about enabling and configuring ISIS:

Navigation Path

- (Device view) Select **Platform > Routing > ISIS** from the Device Policy selector.

- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > ISIS** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Select Enable ISIS to enable Intermediate System-to-Intermediate System routing protocol on the selected ASA device.

About ISIS

Intermediate System-to-Intermediate System (ISIS) routing protocol is a link-state Interior Gateway Protocol (IGP). Link-state protocols are characterized by the propagation of the information required to build a complete network connectivity map on each participating router. That map is then used to calculate the shortest path to destinations. The IOS ISIS implementation supports CLNP, IPv4, and IPv6.

A routing domain may be divided into one or more sub-domains. Each sub-domain is referred to as an area and is assigned an area address. Routing within an area is referred to as Level-1 routing. Routing between Level-1 areas is referred to as Level-2 routing. A router in OSI terminology is referred to as an Intermediate System (IS). An IS may operate at Level 1, Level 2, or both. ISs that operate at Level 1 exchange routing information with other Level-1 ISs in the same area. ISs that operate at Level 2 exchange routing information with other Level-2 routers regardless of whether they are in the same Level-1 area. The set of Level-2 routers and the links that interconnect them form the Level-2 sub-domain, which must not be partitioned in order for routing to work properly.

General Tab

Use the General tab to configure BGP settings such as Best Path Selection, Neighbor Timers, and Graceful Restart.

Navigation Path

You can access the Neighbors tab from the BGP page (see [Configuring BGP](#), on page 2).

Related Topics

- [Configuring BGP](#), on page 2
- [About BGP](#), on page 3
- [IPv4 Family Tab](#), on page 7

Field Reference

Table 30: General Tab

Element	Description
Limit the number of AS numbers in AS_PATH attribute of received routes	Restricts the number of AS numbers in AS_PATH attribute to a specific number. Valid values are from 1 to 254.
Log Neighbor Changes	Enables logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability.

Element	Description
Use TCP path MTU Discovery	Enables the use of the Path MTU Discovery technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation.
Enable fast external failover	Resets the external BGP session immediately upon link failure.
Enforce that the first AS is peer's AS for EBGp routes	Discards incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system.
Use dot notation for AS numbers	Splits the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65535 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation.
Best Path Selection	
Default local preference	Specify a value between 0 and 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers and access servers in the local autonomous system.
Allow comparing MED from different neighbors	Allows the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
Compare Router-id for identical EBGp paths	Compares similar paths received from external BGP peers during the best path selection process and switches the best path to the route with the lowest router ID.
Pick the best MED path among paths advertised from the neighboring AS	Enables MED comparison among paths learned from confederation peers. The comparison between MEDs is made only if no external autonomous systems are there in the path.
Treat missing MED as the least preferred one	Considers the missing MED attribute as having a value of infinity, making the path the least desirable; therefore, a path with a missing MED is least preferred.
Neighbor Timers	
Keepalive Interval	Enter the time interval for which the BGP neighbor remains active after not sending a keepalive message. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.
Hold Time	Enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. The default values is 180 seconds.
Min Hold Time	(Optional) Enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. Specify a value from 0 to 65535.

Element	Description
Graceful Restart (Use in failover or spanned cluster mode) (ASA 9.3.1+ only)	
Enable Graceful Restart	Enables ASA peers to avoid a routing flap following a switchover.
Restart Time	Specify the time duration that ASA peers will wait to delete stale routes before a BGP open message is received. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
Stalepath Time	Enter the time duration that the ASA will wait before deleting stale routes after an end of record (EOR) message is received from the restarting ASA. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.

IPv4 Family Tab

Use the IPv4 Family tab on the BGP page to enable and configure IPv4 settings for BGP.

Navigation Path

You can access the IPv4 Family tab from the BGP page. For more information about the BGP page, see [Configuring BGP](#), on page 2.

Related Topics

- [About BGP](#), on page 3
- [General Tab](#), on page 5

Field Reference

Table 31: IPv4 Family - Aggregate Address Tab

Element	Description
Enable IPv4 Family	Enables configuration of routing sessions that use standard IPv4 address prefixes.
General	Use this panel to configure general IPv4 settings such as Best Path Selection, Neighbor Timers, and Graceful Restart. See IPv4 Family - General Tab , on page 8 for more about these definitions.
Aggregate Address	Use this panel to define the aggregation of specific routes into one route. Specify a value for the aggregate timer (in seconds) in the Aggregate Timer field. Valid values are 0 or any value between 6 and 60. The default value is 30. See Add/Edit Aggregate Address Dialog Box , on page 24 for more about these definitions.

Element	Description
Filtering	Use this panel to filter routes or networks received in incoming BGP updates. See Add/Edit Filter Dialog Box , on page 11 for more about these definitions.
Neighbor	Use this panel to define BGP neighbors and neighbor settings. See Add/Edit Neighbor Dialog Box , on page 12 for more about these definitions.
Networks	Use this panel to define the networks to be advertised by the BGP routing process. See Add/Edit Network Dialog Box , on page 18 for more about these definitions.
Redistribution	Use this panel to define the conditions for redistributing routes from another routing domain into BGP. See Add/Edit Redistribution Dialog Box , on page 19 for more about these definitions.
Route Injection	Use this panel to define the routes to be conditionally injected into the BGP routing table. See Add/Edit Route Injection Dialog Box , on page 20 for more about these definitions.

IPv4 Family Tab—General Tab

Field Reference

Table 32: ISIS IPv4 Family Tab—General Tab

Element	Description
Perform Adjacency Check	Check the ‘Perform adjacency check’ check box for the router to check on nearby IS routers.
Distance	
Administrative Distance	In the Administrative Distance field, enter a distance assigned to routes discovered by IS-IS protocol. Administrative distance is a parameter used to compare routes among different routing protocols. In general, the higher the value, the lower the trust rating. And administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. The range is 1 to 255. The default is 115.
Maximum No. of Forward Paths	Enter the maximum number of IS routes that can be installed in a routing table. The range is 1 to 8, the default is 4.
Distribute Default Route	Check the Distribute default route check box to configure an IS routing process to distribute a default route, and then choose the default route from the Route Map Object selector.
ISIS Metrics	

Element	Description
Global ISIS Metric Level 1	<p>Enter a number specifying the metric.</p> <p>The range depends on the TLV Style that you select. The default is 10.</p> <ul style="list-style-type: none"> • If you select Use old style of TLVs with narrow metric, the range is 1 to 63. • If you select Use new style of TLVs to carry wider metric, the range is 1 to 16777214. • If you select Send and accept both styles of TLVs during transition, the range is 1 to 16777214.
Global ISIS Metric Level 2	<p>Enter a number specifying the metric.</p> <p>The range depends on the TLV Style that you select. The default is 10.</p> <ul style="list-style-type: none"> • If you select Use old style of TLVs with narrow metric, the range is 1 to 63. • If you select Use new style of TLVs to carry wider metric, the range is 1 to 16777214. • If you select Send and accept both styles of TLVs during transition, the range is 1 to 16777214.
TLV Style	<p>Select one of the following Type, Length, and Values:</p> <ul style="list-style-type: none"> • Use old style of TLVs with narrow metric • Use new style of TLVs to carry wider metric • Send and accept both styles of TLVs during transition
Accept both styles of TLVs during transition	<p>If you selected one of the first two options in TLV Style, you can select this option,</p>
Apply metric style to	<p>Select one of the following:</p> <ul style="list-style-type: none"> • Level 1 • Level 2 • Both <p>The default is Level 1.</p>

IPv4 Family Tab—SPF Tab

Field Reference

Table 33: ISIS IPv4 FamilyTab—SPF Tab

Element	Description
Shortest Path First	
Honour external metrics during SPF calculations	Check this check box to have the SPF calculations include external metrics.
Signal other routers to not use this router as an intermediate hop in their SPF calculations	Check this check box if you want to exclude this device, and configure the following:
Specify on-startup behavior	If you select this element you must choose one of the following options: <ul style="list-style-type: none"> • Advertise overself as overloaded until BGP has converged • Specify time to advertise overself as overloaded after reboot—Specify the time in the range of 5 to 86400 seconds.
Don't advertise IP prefixes learned from other protocols when overload bit is set	Check this check box to exclude IP prefixes.
Don't advertise IP prefixes learned from another ISIS level when overload bit is set	Check this check box to exclude IP prefixes.
Minimum interval between partial route calculations	
PRC Interval	Enter an amount of time for the router to wait between partial route calculations (PRCs). The range is 1 to 120 seconds. The default is 5 seconds.
Initial wait for PRC	Enter the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 120,000 milliseconds. The default is 2000 milliseconds.
Minimum wait between first and second PRC	Enter the amount of time in milliseconds that you want the router to wait between PRCs. The range is 1 to 120,000 milliseconds. The default is 5000 milliseconds.
Minimum interval between SPF calculations	
Configure parameters for level 1	
SPF calculation interval	Enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.

Element	Description
Initial wait for SPF calculation	Enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.
Minimum wait between first and second SPF calculation	Enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.
Configure parameters for level 2	
SPF calculation interval	Enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
Initial wait for SPF calculation	Enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.
Minimum wait between first and second SPF calculation	Enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.

IPv4 Family Tab—Redistribution Tab

Use the Add/Edit button to add a new Redistribution route or edit an existing row.

Field Reference

Table 34: ISIS IPv4 Family Tab—Redistribution Tab

Element	Description
Source Protocol	From the Source Protocol drop-down list, choose the protocol (BGP, Connected, EIGRP, OSPF, RIP, or Static) from which you want to redistribute routes into the ISIS domain.
Process ID	Enter the Process ID for the source protocol.
Route Level	From the Route Level drop-down list, choose Level-1, Level- 2, or Level 1-2.
Metric	In the Metric field, enter a metric for the redistributed route. The range is 1 to 4294967295.
Metric Type	For the Metric Type, click the internal or external radio button.
ISIS Inter Area Route Levels	
Source ISIS Level	Select Level 1 or Level 2. The default is Level 1.
Destination ISIS Level	Select Level 1 or Level 2. The default is Level 1.
Distribution List	Select from the available Access Control List or add new.

Element	Description
Route Map	Choose a route map from the Route Map Object selector that should be examined to filter the networks to be redistributed, or click Add to add a new route map or edit an existing route map.
Match	Check one or more of the Match check boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

IPv6 Family Tab

Use the IPv6 Family tab on the BGP page to enable and configure IPv6 settings for BGP.

Navigation Path

You can access the IPv6 Family tab from the BGP page. For more information about the BGP page, see [Configuring BGP](#), on page 2.

Related Topics

- [About BGP](#), on page 3
- [General Tab](#), on page 5

Field Reference

Table 35: IPv6 Family - Aggregate Address Tab

Element	Description
Enable IPv6 Family	Enables configuration of routing sessions that use standard IPv6 address prefixes.
General	Use this panel to configure general IPv6 settings. See IPv6 Family - General Tab , on page 22 for more about these definitions.
Aggregate Address	Use this panel to define the aggregation of specific routes into one route. Specify a value for the aggregate timer (in seconds) in the Aggregate Timer field. Valid values are 0 or any value between 6 and 60. The default value is 30. See Add/Edit Aggregate Address Dialog Box , on page 24 for more about these definitions.
Neighbor	Use this panel to define BGP neighbors and neighbor settings. See Add/Edit Neighbor Dialog Box , on page 25 for more about these definitions.
Networks	Use this panel to define the networks to be advertised by the BGP routing process. See Add/Edit Network Dialog Box , on page 31 for more about these definitions.
Redistribution	Use this panel to define the conditions for redistributing routes from another routing domain into BGP. See Add/Edit Redistribution Dialog Box , on page 32 for more about these definitions.

Element	Description
Route Injection	Use this panel to define the routes to be conditionally injected into the BGP routing table. See Add/Edit Route Injection Dialog Box , on page 33 for more about these definitions.

IPv6 Family Tab—General Tab

Field Reference

Table 36: ISIS IPv6 Family Tab—General Tab

Element	Description
Perform Adjacency Check	Check the ‘Perform adjacency check’ check box for the router to check on nearby IS routers.
Distance	
Administrative Distance	In the Administrative Distance field, enter a distance assigned to routes discovered by ISIS protocol. Administrative distance is a parameter used to compare routes among different routing protocols. In general, the higher the value, the lower the trust rating. And administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. The range is 1 to 255. The default is 115.
Maximum No. of Forward Paths	Enter the maximum number of IS routes that can be installed in a routing table. The range is 1 to 8. The default is 4.
Distribute Default Route	Check the Distribute default route check box to configure an IS routing process to distribute a default route, and then choose the default route from the Route Map Object selector.

IPv6 Family Tab—SPF Tab

Field Reference

Table 37: ISIS IPv6 FamilyTab—SPF Tab

Element	Description
Shortest Path First	
Signal other routers to not use this router as an intermediate hop in their SPF calculations	Check this check box if you want to exclude this device, and configure the following:

Element	Description
Specify on-startup behavior	If you select this element you must choose one of the following options: <ul style="list-style-type: none"> • Advertise overself as overloaded until BGP has converged • Specify time to advertise overself as overloaded after reboot—Specify the time in the range of 5 to 86400 seconds.
Don't advertise IP prefixes learned from other protocols when overload bit is set	Check this check box to exclude IP prefixes.
Don't advertise IP prefixes learned from another ISIS level when overload bit is set	Check this check box to exclude IP prefixes.
Minimum interval between partial route calculations	
PRC Interval	Enter an amount of time for the router to wait between partial route calculations (PRCs). The range is 1 to 120 seconds. The default is 5 seconds.
Initial wait for PRC	Enter the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 120.000 milliseconds. The default is 2000 milliseconds.
Minimum wait between first and second PRC	Enter the amount of time in milliseconds that you want the router to wait between PRCs. The range is 1 to 120,000 milliseconds. The default is 5000 milliseconds.
Minimum interval between SPF calculations	
Configure parameters for level 1	
SPF calculation interval	Enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
Initial wait for SPF calculation	Enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
Minimum wait between first and second SPF calculation	Enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.
Configure parameters for level 2	
SPF calculation interval	Enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.

Element	Description
Initial wait for SPF calculation	Enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.
Minimum wait between first and second SPF calculation	Enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.

IPv6 Family Tab—Redistribution Tab

Use the Add/Edit button to add or edit Redistribution routes.

Field Reference

Table 38: ISIS IPv6 Family Tab—Redistribution Tab

Element	Description
Source Protocol	From the Source Protocol drop-down list, choose the protocol (BGP, Connected, EIGRP, OSPF, RIP, or Static) from which you want to redistribute routes into the ISIS domain.
Process ID	Enter the Process ID for the source protocol.
Route Level	From the Route Level drop-down list, choose Level-1, Level- 2, or Level 1-2.
Metric	In the Metric field, enter a metric for the redistributed route. The range is 1 to 4294967295.
Metric Type	For the Metric Type, click the internal or external radio button.
ISIS Inter Area Route Levels	
Source ISIS Level	Select Level 1 or Level 2. The default is Level 1.
Destination ISIS Level	Select Level 1 or Level 2. The default is Level 1.
Distribution List	Select from the available Access Control List or add new.
Route Map	Choose a route map from the Route Map Object selector that should be examined to filter the networks to be redistributed, or click Add to add a new route map or edit an existing route map.
Match	Check one or more of the Match check boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

IPv6 Family Tab—Summary Prefix

You must configure at least one Network Entity Title entry to proceed.

See [Network Entity Title Tab](#) , on page 65 for more information.

Use the Add/Edit button to add or edit Summary Prefix.

Field Reference

Table 39: ISIS IPv6 Family Tab—Summary Prefix Tab

Element	Description
IPv6 Summary Prefix	IPv6 prefix in the form X.X.X.X::X/0-128
Apply Summary Prefix into	<p>Select Level 1, Level 2, or Both.</p> <p>Level 1: Only routes redistributed into Level 1 are summarized with the configured address and mask value.</p> <p>Level 2: Routes learned by Level 1 routing are summarized into the Level 2 backbone with the configured address and mask value. Redistributed routes into Level 2 ISIS are also summarized.</p> <p>Both: Summary routes are applied when redistributing routes into Level 1 and Level2 ISIS and when Level 2 ISIS advertises Level 1 routes as reachable in it area.</p>

Authentication Tab

Field Reference

Table 40: ISIS Authentication Tab

Element	Description
Configure authentication parameter for level 1	
Type	Select a Type from the drop-down list.
Key	Enter the key to authenticate ISIS updates. The key can include up to 16 characters.
Confirm	Confirm the key.
Send only	Click Enable or Disable depending on whether you want Send Only enabled.
Mode	Choose the authentication mode by clicking either the Disabled, MD5, or Clear Text radio buttons.
Area password	Enter the Area password and confirm the same in the next textbox.
Configure authentication parameter for level 2	
Type	Select a Type from the drop-down list.
Key	Enter the key to authenticate ISIS updates. The key can include up to 16 characters.

Element	Description
Confirm	Confirm the key.
Send only	Click Enable or Disable depending on whether you want Send Only enabled.
Mode	Choose the authentication mode by clicking either the Disabled, MD5, or Clear Text radio buttons.
Domain password	Enter the Domain password and confirm the same.

Link State Packet Tab

Field Reference

Table 41: ISIS Link State Packet Tab

Element	Description
Ignore LSP Errors	Check the Ignore LSP Errors check box to allow the ASA to ignore LSP packets that are received with internal checksum errors rather than purging the LSPs.
Flood LSPs before running SPF	Check this box to fast-flood and fill LSPs before running SPF. If you select this option, enter the number of LSPs to be flooded in the range of 1 to 15. This parameter sends a specified number of LSPs from the ASA. If no LSP number is specified, the default of 5 is used. The LSPs invoke SPF before running SPF. Cisco recommends that you enable fast flooding, because then you speed up the LSP flooding process, which improves overall network convergence time. The default value is 5.
Suppress IP prefixes	To suppress IP prefixes, check the Suppress IP prefixes check box, and then check one of the following. In networks where there is no limit placed on the number of redistributed routes into IS-IS, it is possible that the LSP can become full and routes will be dropped. Use these options to control which routes are suppressed when the PDU becomes full.
Don't advertise IP prefixes learned from another ISIS level when ran out of LSP fragments	Suppresses any routes coming from another level. For example, if the Level-2 LSP becomes full, routes from Level 1 are suppressed.
Don't advertise IP prefixes learned from other protocols when ran out of LSP fragments	Suppresses any redistributed routes on the ASA.
LSP General Interval	
LSP Interval Parameters for level 1	

Element	Description
LSP Calculation Interval	<p>Enter the interval of time in seconds between transmission of each LSP. The range is 1-120 seconds. The default is 5.</p> <p>The number should be greater than the expected round-trip delay between any two ASAs on the attached network. The number should be conservative or needless transmission results. Retransmissions occur only when LSPs are dropped. So setting the number to a higher value has little effect on reconvergence. The more neighbors the ASAs have, and the more paths over which LSPs can be flooded, the higher you can make this value.</p>
Initial wait for LSP calculation	Enter the time in milliseconds specifying the initial wait time before the first LSP is generated. The range is 1 to 120,000. The default is 50.
Minimum wait between first and second	Enter the time in milliseconds between the first and second LSP generation. The range is 1 to 120,000. The default is 5000.
LSP Interval Parameters for level 2	
Use level 1 parameter also for level 2	If you want the values you configured for Level 1 to also apply to Level 2, check the Use level 1 parameters also for level 2 check box.
LSP Calculation Interval	<p>Enter the interval of time in seconds between transmission of each LSP. The range is 1-120 seconds. The default is 5.</p> <p>The number should be greater than the expected round-trip delay between any two ASAs on the attached network. The number should be conservative or needless transmission results. Retransmissions occur only when LSPs are dropped. So setting the number to a higher value has little effect on reconvergence. The more neighbors the ASAs have, and the more paths over which LSPs can be flooded, the higher you can make this value.</p>
Initial wait for LSP calculation	Enter the time in milliseconds specifying the initial wait time before the first LSP is generated. The range is 1 to 120,000. The default is 50.
Minimum wait between first and second	Enter the time in milliseconds between the first and second LSP generation. The range is 1 to 120,000. The default is 5000.
Maximum LSP size	In the Maximum LSP size field, enter the number of seconds. The range is 128 to 4352. The default is 1492.
LSP refresh interval	<p>In the LSP refresh interval field, enter the number of seconds at which LSPs are refreshed. The range is 1 to 655535. The default is 900.</p> <p>The refresh interval determines the rate at which the software periodically transmits in LSPs the route topology information that it originates. This is done to keep the database information from becoming too old.</p> <p>Reducing the refresh interval reduces the amount of time that undetected link state database corruption can persist at the cost of increased link utilization. (This is an extremely unlikely event, however, because there are other safeguards against corruption.) Increasing the interval reduces the link utilization caused by the flooding of refreshed packets (although this utilization is very small).</p>

Element	Description
Maximum LSP lifetime	<p>In the Maximum LSP lifetime field, enter the maximum number of seconds that LSPs can remain in a router's database without being refreshed. The range is 1 to 65535. The default is 1200 (20 minutes).</p> <p>You might need to adjust this parameter if you change the LSP refresh interval. LSPs must be periodically refreshed before their lifetimes expire. The value set for LSP refresh interval should be less than the value set for the maximum LSP lifetime; otherwise LSPs will time out before they are refreshed. If you make the LSP lifetime too low compared to the LSP refresh interval, the LSP refresh interval is automatically reduced to prevent the LSPs from timing out.</p>

Summary Address Tab

Use the Add/Edit button to add or edit Summary Addresses.

Field Reference

Table 42: ISIS Summary Address Tab

Element	Description
IP address	Enter the IP address of the summary route.
Net Mask	Choose or enter the network mask to apply to the IP address.
Select level	Select the Level 1, Level 2, or Level 1 and 2 radio button depending on which levels you want to receive summary addresses.
Tag	In the Tag field, enter a number for the tag. The range is from 1 to 4294967295.
Metric	In the Metric field, enter the metric that will be applied to the summary route. The range is from 1 to 4294967295. The default value is 10.

Network Entity Title Tab

Use the Add/Edit button to add to edit Network Entity Title.

Field Reference

Table 43: ISIS Network Entity Title Tab

Element	Description
Network Entity Title (NET)	Enter a value in the address format 48.0000.1111.2222.00. The total length of NET address must be between 16 and 40 characters.

Element	Description
NET Pool	<p>Click Select to open the NET Pool Object Selector dialog box. You can add and edit NET Pool Objects using this dialog box. For more information about how to add or edit NET Pool Objects, see Add or Edit NET Pool Object Dialog Box.</p> <p>The NET Pool is applicable only for cluster devices in individual mode.</p> <p>Network Entity Title (NET) is not applicable for cluster devices in individual mode.</p>
Maximum allowed NET	Enter a NET value in the range of 3 to 254. The default value is 3.

Interface Tab

Use the Interface tab to configure interface-specific OSPF authentication routing properties.

Navigation Path

You can access the Interface tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#), on page 80.

Related Topics

- [Add/Edit Interface Dialog Box](#), on page 106

Field Reference

Table 44: Interface Tab

Element	Description
Interface	The name of the interface to which the configuration applies.
Authentication	<p>The type of OSPF authentication enabled on the interface. The authentication type can be one of the following values:</p> <ul style="list-style-type: none"> • None—OSPF authentication is disabled. • Password—Clear text password authentication is enabled. • MD5—MD5 authentication is enabled. • Area—The authentication type specified for the area is enabled on the interface. Area authentication is the default value for interfaces. However, area authentication is disabled by default. So, unless you previously specified an area authentication type, interfaces showing Area authentication have authentication disabled. • Key Chain—Key chain authentication is enabled.
Point-to-Point	Displays “true” if the interface is set to non-broadcast (point-to-point). Displays “false” if the interface is set to broadcast.

Element	Description
Cost	The cost of sending a packet through the interface.
Priority	The OSPF priority assigned to the interface.
MTU Ignore	Displays “false” if MTU mismatch detection is enabled. Displays “true” if the MTU mismatch detection is disabled.
Database Filter	Displays “true” if outgoing LSAs are filtered during synchronization and flooding. Displays “false” if filtering is not enabled.
Hello Interval	The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds. The default value is 10 seconds.
Transmit Delay	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds. The default value is 1 second.
Retransmit Interval	The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Dead Interval	The interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535. The default value of this setting is four times the interval set by the Hello Interval field.
Hello Multiplier (ASA 9.2(1)+ only)	The number of hello packets to be sent per second. Valid values are between 3 and 20.

Interface Tab—General Tab

Field Reference

Table 45: ISIS Interfaces Tab—General Tab

Element	Description
Interface	Select the interface from available interfaces.

Element	Description
Shutdown ISIS on this interface	Shutdown ISIS on this interface—Lets you disable the IS-IS protocol for this interface without removing the configuration parameters. The IS-IS protocol will not form any adjacencies on this interface and the IP address of this interface will be put into the LSP that is generated by the ASA.
Enable ISIS on this interface	Enables IS-IS protocol on the selected interface.
Enable IPv6 ISIS on this interface	Enables IPv6 IS-IS routing on the selected interface.
Priority for level 1	Lets you set a priority for Level 1. The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS. The range is 0 to 127. The default is 64.
Priority for level 2	Lets you set a priority for Level 2. The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS. The range is 0 to 127. The default is 64.
Tag	Sets a tag on the IP address configured for an interface when this IP prefix is put into an ISIS LSP.
CSNP Interval for level 1	Sets the Complete Sequence Number PDUs (CSNPs) interval in seconds between transmission of CSNPs on multiaccess networks for Level 1. This interval only applies for the designated router. The range is from 0 to 65535. The default is 10 seconds.
CSNP Interval for level 2	Sets the Complete Sequence Number PDUs (CSNPs) interval in seconds between transmission of CSNPs on multiaccess networks for Level 2. This interval only applies for the designated router. The range is from 0 to 65535. The default is 10 seconds.
Adjacency filter	Filters the establishment of IS-IS adjacencies.
Match all area addresses	All NSAP addresses must match the filter to accept the adjacency. If not specified (the default), only one address must match the filter for the adjacency to be accepted.

Interface Tab—Authentication Tab

Field Reference

Table 46: ISIS Interfaces Tab—Authentication Tab

Element	Description
Level 1 parameters	

Element	Description
Key Type	Select Clear Text or Encrypted.
Key	Enter the key to authenticate IS-IS updates. The range is 0 to 8 characters. If no password is configured with the Key option, no key authentication is performed. Note If you selected Key Type as Clear Text, you can enter a maximum of 17 characters in the Key field. If you selected Key Type as Encrypted, you can enter a maximum of 50 characters in the Key field.
Send only	For Send only click the Enable or Disable radio button. Choosing Send only causes the system only to insert the password into the SNPs, but not check the password in SNPs that it receives. Use this keyword during a software upgrade to ease the transition. The default is disabled.
Mode	Choose the authentication mode by checking the Mode check box and then choosing MD5 or Text from the drop-down list.
Password	Enter a password. Note You can select either Mode or enter a password value.
Level 2 parameters	
Key Type	Select Clear Text or Encrypted.
Key	Enter the key to authenticate IS-IS updates. The range is 0 to 8 characters. If no password is configured with the Key option, no key authentication is performed. Note If you selected Key Type as Clear Text, you can enter a maximum of 17 characters in the Key field. If you selected Key Type as Encrypted, you can enter a maximum of 50 characters in the Key field.
Send only	For Send only click the Enable or Disable radio button. Choosing Send only causes the system only to insert the password into the SNPs, but not check the password in SNPs that it receives. Use this keyword during a software upgrade to ease the transition. The default is disabled.
Mode	Choose the authentication mode by checking the Mode check box and then choosing MD5 or Text from the drop-down list.
Password	Enter a password. Note You can select either Mode or enter a password value.

Interface Tab—Hello Padding Tab

Field Reference

Table 47: ISIS Interfaces Tab—Hello Padding Tab

Element	Description
Hello Padding	Enables Hello Padding. IS-IS hellos are padded to the full maximum transmission unit (MTU) size. Padding IS-IS hellos to the full MTU allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces.
Minimal holdtime 1 second for level 1	Enables the holdtime (in seconds) that the LSP remains valid for Level 1.
Hello interval for level 1	Specifies the length of time in seconds between hello packets for Level 1. The range is 1 to 65535. The default is 10.
Minimal holdtime 1 second for level 2	Enables the holdtime (in seconds) that the LSP remains valid for Level 2.
Hello interval for level 2	Specifies the length of time in seconds between hello packets for Level 2. The range is 1 to 65535. The default is 10.
Hello multiplier for level 1	Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency is down for Level 1. The advertised hold time in IS-IS hello packets will be set to the hello multiplier times the hello interval. Neighbors will declare an adjacency to this router down after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on a per-interface basis, and can be different between different routers in one area. The range is 3 to 1000. The default is 3.
Hello multiplier for level 2	Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency is down for Level 2. The advertised hold time in IS-IS hello packets will be set to the hello multiplier times the hello interval. Neighbors will declare an adjacency to this router down after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on a per-interface basis, and can be different between different routers in one area. The range is 3 to 1000. The default is 3.
Configure Circuit Type	Specifies whether the interface is configured for local routing (level 1), area routing (Level 2), or both local and area routing (Level 1-2).

Interface Tab—LSP Settings Tab

Field Reference

Table 48: ISIS Interfaces Tab—LSP Settings Tab

Element	Description
Advertise ISIS Prefix	<p>Allows the advertising of IP prefixes of connected networks in the LSP advertisements per IS-IS interface.</p> <p>Disabling this option is an IS-IS mechanism to exclude IP prefixed of connected network from LSP advertisements thereby reducing IS-IS convergence time.</p>
Retransmit Interval	<p>Specifies the amount of time in seconds between retransmission of each IS-IS LSP on a point-to-point link.</p> <p>The number should be greater than the expected round-trip delay between any two routers on the attached network. The range is 0 to 65535. The default is 5.</p>
Retransmit Throttle Interval	<p>Specifies the amount of time in milliseconds between retransmissions on each IS-IS LSP on a point-to-point interface.</p> <p>This option may be useful in very large networks with many LSPs and many interfaces as a way of controlling LSP retransmission traffic. This option controls the rate at which LSPs can be re-sent on the interface. The range is 0 to 65535. The default is 33.</p>
LSP Interval	<p>Specifies the time delay in millisecond between successive IS-IS LSP transmissions.</p> <p>In topologies with a large number of IS-IS neighbors and interfaces, a router may have difficulty with the CPU load imposed by LSP transmission and reception. This option allows the LSP transmission rate (and by implication the reception rate of other systems) to be reduced. The range is 1 to 4294967295. The default is 33.</p>

Interface Tab—Metrics Tab

Field Reference

Table 49: ISIS Interfaces Tab—Metrics Tab

Element	Description
Metrics for level 1	
Use maximum metric value	Specifies the metric assigned to the link and used to calculate the cost from each other router via the links in the network to other destinations. This is enabled by default.
Default metric	Enter the number for the metric. The range is 1 to 16777214.
Metrics for level 2	

Element	Description
Use maximum metric value	Specifies the metric assigned to the link and used to calculate the cost from each other router via the links in the network to other destinations. This is enabled by default.
Default metric	Enter the number for the metric. The range is 1 to 16777214.

Passive Interfaces Tab

The Passive Interfaces tab enables you to allow or suppress routing updates on an interface. Only interfaces configured with a name can be suppressed from sending routing updates.

Field Reference

Table 50: ISIS Network Entity Title Tab

Element	Description
Passive Interface	<p>Select from the following options:</p> <ul style="list-style-type: none"> • None—No Interface is selected. • Default—Open the Interfaces Selector dialog to select interfaces that you want to exclude. By default all interfaces are selected. • Specified Interfaces—Open the Interfaces Selector dialog to select interfaces that you want to select and include.

Configuring BFD Routing

The BFD page provides two tabs for configuring BFD (Bidirectional Forwarding Detection) routing on a firewall device. The following topics provide detailed information on configuring BFD.

Navigation Path

- (Device view) Select **Platform > Routing > BFD** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > BFD** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- [About BFD, on page 73](#)
- [Create BFD Template, on page 76](#)
- [Add/Edit BFD Map Dialog Box, on page 78](#)
- [Add/ Edit BFD Interface Dialog Box, on page 79](#)

About BFD

Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. BFD operates in a unicast, point-to-point mode on top of any data protocol being forwarded between two systems. Packets are carried in the payload of the encapsulating protocol appropriate for the media and the network.

BFD provides a consistent failure detection method for network administrators in addition to fast forwarding path failure detection. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning are easier and reconvergence time is consistent and predictable.

BFD Asynchronous Mode and Echo Function

BFD can operate in asynchronous mode with or without the echo function enabled.

Asynchronous Mode

In asynchronous mode, the systems periodically send BFD control packets to one another, and if a number of those packets in a row are not received by the other system, the session is declared to be down. Pure asynchronous mode (without the Echo function) is useful because it requires half as many packets to achieve a particular detection time as the Echo function requires.

BFD Echo Function

The BFD echo function sends echo packets from the forwarding engine to the directly-connected single-hop BFD neighbor. The echo packets are sent by the forwarding engine and forwarded back along the same path to perform detection. The BFD session at the other end does not participate in the actual forwarding of the echo packets. Because the echo function and the forwarding engine are responsible for the detection process, the number of BFD control packets that are sent out between BFD neighbors is reduced. And also because the forwarding engine is testing the forwarding path on the remote neighbor system without involving the remote system, the inter-packet delay variance is improved. This results in quicker failure detection times.

When the echo function is enabled, BFD can use the slow timer to slow down the asynchronous session and reduce the number of BFD control packets that are sent between BFD neighbors, which reduces processing overhead while at the same time delivering faster failure detection.




Note The echo function is not supported for IPv4 multi-hop or IPv6 single-hop BFD neighbors.

You can enable BFD at the interface and routing protocol levels. You must configure BFD on both systems (BFD peers). After you enable BFD on the interfaces and at the router level for the appropriate routing protocols, a BFD session is created, BFD timers are negotiated, and the BFD peers begin to send BFD control packets to each other at the negotiated level.

BFD Session Establishment

The following example shows the ASA and a neighboring router running Border Gateway Protocol (BGP). At the time when both devices come up, there is no BFD session established between them.

Figure 1: BFD Session Initiated



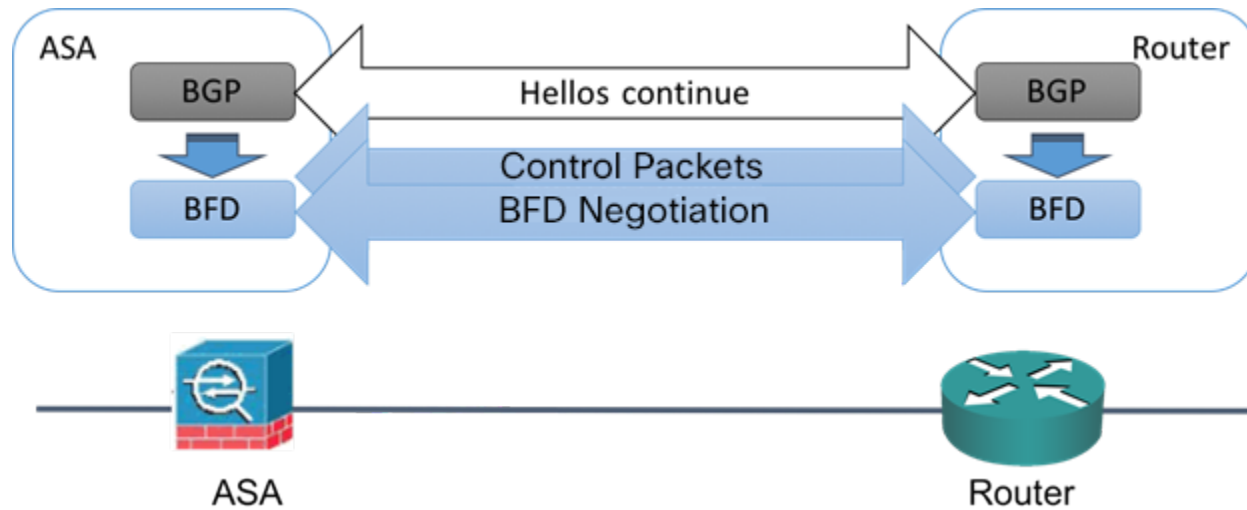
After BGP identifies its BGP neighbor, it bootstraps the BFD process with the IP address of the neighbor. BFD does not discover its peers dynamically. It relies on the configured routing protocols to tell it which IP addresses to use and which peer relationships to form.

The BFD on the router and the BFD on the ASA form a BFD control packet and start sending the packets to each other at a one-second interval until the BFD session is established. The initial control packets from either system are very similar, for example, the Vers, Diag, H, D, P, and F bits are all set to zero, and the State is set to Down. The My Discriminator field is set to a value that is unique on the transmitting device. The Your Discriminator field is set to zero because the BFD session has not yet been established. The TX and RX timers are set to the values found in the configuration of the device.

After the remote BFD device receives a BFD control packet during the session initiation phase, it copies the value of the My Discriminator field into its own Your Discriminator field and the transition from Down state to Init state and then eventually to Up state occurs. Once both systems see their own Discriminators in each other's control packets, the session is officially established.

The following illustration shows the established BFD connection.

Figure 2: BFD Session Established



BFD Timer Negotiation

BFD devices must negotiate the BFD timers to control and synchronize the send rate of BFD control packets.

A device needs to ensure the following before it can negotiate a BFD timer:

- That its peer device saw the packet containing the proposed timers of the local device
- That it never sends BFD control packets faster than the peer is configured to receive them
- That the peer never sends BFD control packets faster than the local system is configured to receive them

The setting of the Your Discriminator field and the H bit are sufficient to let the local device that the remote device has seen its packets during the initial timer exchange. After receiving a BFD control packet, each system takes the Required Min RX Interval and compares it to its own Desired Min TX Interval, and then takes the greater (slower) of the two values and uses it as the transmission rate for its BFD packets. The slower of the two systems determines the transmission rate.

When these timers have been negotiated, they can be renegotiated at any time during the session without causing a session reset. The device that changes its timers sets the P bit on all subsequent BFD control packets until it receives a BFD control packet with the F bit set from the remote system. This exchange of bits guards against packets that might otherwise be lost in transit.



Note The setting of the F bit by the remote system does not mean that it accepts the newly proposed timers. It indicates that the remote system has seen the packets in which the timers were changed.

BFD Failure Detection

When the BFD session and timers have been negotiated, the BFD peers send BFD control packets to each other at the negotiated interval. These control packets act as a heartbeat that is very similar to IGP Hello protocol except that the rate is more accelerated.

As long as each BFD peer receives a BFD control packet within the configured detection interval (Required Minimum RX Interval), the BFD session stays up and any routing protocol associated with BFD maintains its adjacencies. If a BFD peer does not receive a control packet within this interval, it informs any clients participating in that BFD session about the failure. The routing protocol determines the appropriate response to that information. The typical response is to terminate the routing protocol peering session and reconverge and thus bypass a failed peer.

Each time a BFD peer successfully receives a BFD control packet in a BFD session, the detection timer for that session is reset to zero. Thus the failure detection is dependent on received packets and NOT when the receiver last transmitted a packet.

BFD Deployment Scenarios

The following describes how BFD operates in these specific scenarios.

Failover

In a failover scenario, BFD sessions are established and maintained between the active unit and the neighbor unit. Standby units do not maintain any BFD sessions with the neighbors. When a failover happens, the new active unit must initiate session establishment with the neighbor because session information is not synched between active and standby units.

For a graceful restart/NSF scenario, the client (BGP IPv4/IPv6) is responsible for notifying its neighbor about the event. When the neighbor receives the information, it keeps the RIB table until failover is complete. During failover, the BFD and the BGP sessions go down on the device. When the failover is complete, a new BFD session between the neighbors is established when the BGP session comes up.

Spanned EtherChannel and L2 Cluster

In a Spanned EtherChannel cluster scenario, the BFD session is established and maintained between the primary unit and its neighbor. Subordinate units do not maintain any BFD sessions with the neighbors. If a BFD packet is routed to the subordinate unit because of load balancing on the switch, the subordinate unit must forward this packet to the primary unit through the cluster link. When a cluster switchover happens, the new primary unit must initiate session establishment with the neighbor because session information is not synched between primary and subordinate units.

Individual Interface Mode and L3 Cluster

In an individual interface mode cluster scenario, individual units maintain their BFD sessions with their neighbors.

Create BFD Template

This section describes the steps required to create a BFD template policy object. The BFD template specifies a set of BFD interval values. BFD interval values as configured in the BFD template are not specific to a single interface. You can also configure authentication for single-hop and multi-hop sessions. You can enable Echo on single-hop only.

Navigation Path

Select **Manage > Policy Objects**, then select **BFD Template** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Field Reference

Table 51: Add/Edit BFD Template

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects .
Description	An optional description of the object.
Config Mode	Specify if there is a single IP hop or multiple IP hops between a BFD source and destination associated with an interface.
Enable Echo	(Optional) Select to enable echo. When enabled, echo packets are sent by the forwarding engine and forwarded back along the same path in order to perform detection. Note This is applicable only for single hop configuration mode.
Interval tab (Optional)	
Interval Type	Specify if you want to define the interval type in microseconds or milliseconds. The default interval type is None.
Transmit and Receive Values Interval Values in Microseconds	This section is enabled, if the Interval type is microseconds. Valid values are between 50000 and 999000 microseconds. Minimum Transmit Value - Enter the minimum transmit interval capability in microseconds. Minimum Receive Value - Enter the minimum receive interval capability in microseconds.
Transmit and Receive Values Interval Values in Milliseconds	This section is enabled, if the Interval type is milliseconds. Valid values are between 50 and 999 milliseconds. Minimum Transmit Value - Enter the minimum transmit interval capability in milliseconds. Minimum Receive Value - Enter the minimum receive interval capability in milliseconds.
Multiplier Value	Enter the number of consecutive BFD control packets that must be missed before BFD declares that a peer is unavailable. The default value is 3. Valid values are between 3 to 50.
Authentication tab (Optional)	
Authentication Type	Select to configure authentication for the BFD template and specify if you want to use an encrypted password or an unencrypted password, for the authentication.

Element	Description
Key Value	Enter a BFD password and confirm it. <ul style="list-style-type: none"> For encrypted BFD templates, the length of the Key value is between 17 and 66 characters. For unencrypted BFD templates of sha-1 or meticulous-sha-1 authentication type the length of the Key Value must be less than 29 characters. For unencrypted BFD templates of md5 or meticulous-md5 authentication type the length of the Key Value must be less than 25 characters.
Key ID	Enter an authentication key ID. This is a shared key ID, that matches the key string.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects .
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden and Understanding Policy Object Overrides for Individual Devices . If you allow device overrides, you can click the Edit button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.

Add/Edit BFD Map Dialog Box

The Add/ Edit BFD Map dialog box lets you create a BFD map containing destinations that you can associate with a multi-hop template. You must have a multi-hop BFD template already configured. For more information see, [Create BFD Template, on page 76](#).

Navigation Path

You can access the Add/ Edit BFD Map dialog box from the Maps tab on the BFD page. Click the Add Row button to add a new BFD map; select an existing BFD map and click the Edit Row button to edit that map.

Related Topics

- [Create BFD Template, on page 76](#)

Field Reference

Table 52: BFD Maps Tab

Element	Description
BFD Template	Select a multi-hop BFD template or add a multi-hop BFD template. For more information see, Create BFD Template, on page 76 .

Element	Description
IP version	Select the appropriate address format for the source and destination - IPv4 or IPv6.
IPv4 Destination/Prefix, IPv4 Source/Prefix	Enter the IPv4 address for the destination and source in the appropriate fields in the x.x.x.x/prefix format.
IPv6 Destination/ Prefix, IPv6 Source/prefix	Enter the IPv6 address for the destination and source in the appropriate fields in the x:x:x:x:x:x/prefix format.
Slow Timers	This reduces the number of BFD control packets, that are sent between BFD neighbors. This slows down the asynchronous session, reduces the processing overhead and results in faster failure detection. The default value for slow timers is 1000 and valid values are between 1000 - 30000.

Add/ Edit BFD Interface Dialog Box

The Add/ Edit BFD Interfaces dialog box lets you bind a BFD template to an interface, configure the baseline BFD session parameters per interface, and enable echo mode per interface.

Navigation Path

You can access the Add/ Edit BFD Interface dialog box from the Interface tab on the BFD page. Click the Add Row button to add a new BFD interface; select an existing BFD interface and click the Edit Row button to edit that map.

Related Topics

- [Create BFD Template, on page 76](#)

Field Reference

Table 53: BFD Interface Tab

Element	Description
Interface	Enter an interface name, select an interface or add an interface role.
BFD Configuration	Select BFD template to select an existing single-hop BFD template or add a single-hop BFD template. Alternately, select BFD interval. For more information see, Create BFD Template, on page 76 .
BFD Interval	
Minimum Transmit Value	Enter the minimum transmit interval capability in milliseconds. Valid values are between 50 and 999 milliseconds
Minimum Receive Value	Enter the minimum receive interval capability milliseconds. Valid values are between 50 and 999 milliseconds

Element	Description
Multiplier	Enter the number of consecutive BFD control packets that must be missed before BFD declares that a peer is unavailable. The default value is 3. Valid values are between 3 to 50.
Echo	(Optional) Select to enable echo. When enabled, echo packets are sent by the forwarding engine and forwarded back along the same path in order to perform detection.

Configuring OSPF

The OSPF page provides ten tabbed panels for configuring OSPF (Open Shortest Path First) routing on a firewall device. The following topics provide detailed information about enabling and configuring OSPF:



Note Depending on the device version that you are configuring, some tabs might not be available.



Note Beginning with ASA version 9.2(1), certain OSPF settings have changed. If you configure a shared policy that uses settings specific to ASA 9.2(1)+, you will receive a validation error if that policy is assigned to a device whose version is earlier than 9.2(1). Likewise, if you configure a shared policy that uses settings that no longer apply to ASA 9.2(1)+, you will receive a validation error if that policy is assigned to an 9.2(1)+ device.

- [About OSPF](#) , on page 81
- [General Tab](#) , on page 5
- [Area Tab](#) , on page 88
- [Range Tab](#) , on page 91
- [Neighbors Tab](#) , on page 42
- [Redistribution Tab](#) , on page 44
- [Virtual Link Tab](#) , on page 96
- [Filtering Tab](#) , on page 99
- [Filter Rule Tab](#) , on page 101
- [Summary Address Tab](#) , on page 47
- [Interface Tab](#) , on page 66
- [Configuring Key Chain](#) , on page 108

Navigation Path

- (Device view) Select **Platform > Routing > OSPF** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > OSPF** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

About OSPF

Open Shortest Path First (OSPF) is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. OSPF propagates link-state advertisements (LSAs) rather than routing table updates. Because only LSAs are exchanged, rather than entire routing tables, OSPF networks converge more quickly than RIP networks.

OSPF supports MD5 and clear-text neighbor authentication. Authentication should be used with all routing protocols whenever possible, because route redistribution between OSPF and other protocols (like RIP) can potentially be used by attackers to subvert routing information.

If NAT is used when OSPF is operating on public and private areas, and if address filtering is required, you need to run two OSPF processes—one process for the public areas and one for the private areas.

A router that has interfaces in multiple areas is called an Area Border Router (ABR). A router that acts as a gateway to redistribute traffic between routers using OSPF and routers using other routing protocols is called an Autonomous System Boundary Router (ASBR).

An ABR uses LSAs to send information about available routes to other OSPF routers. Using ABR type 3 LSA filtering, you can have separate private and public areas with the security appliance acting as an ABR. Type 3 LSAs (inter-area routes) can be filtered from one area to other. This lets you use NAT and OSPF together without advertising private networks.



Note Only type 3 LSAs can be filtered. If you configure the security appliance as an ASBR in a private network, it will send type 5 LSAs describing private networks, which will be broadcast to the entire autonomous system (AS) including public areas.

If NAT is employed but OSPF is only running in public areas, routes to public networks can be redistributed inside the private network, either as default or type 5 AS External LSAs. However, you need to configure static routes for the private networks protected by the security appliance. Also, you should not mix public and private networks on the same security appliance interface.

Related Topics

- [Configuring OSPF](#) , on page 80

General Tab

Use the General panel on the OSPF page to enable up to two OSPF process instances. Each OSPF process has its own associated areas and networks.



Note You cannot enable OSPF if you have RIP enabled.

Navigation Path

You can access the General tab from the OSPF page; see [Configuring OSPF](#), on page 80 for more information.

Related Topics

- [Area Tab](#), on page 88
- [Range Tab](#), on page 91
- [Neighbors Tab](#), on page 42
- [Redistribution Tab](#), on page 44
- [Virtual Link Tab](#), on page 96
- [Filtering Tab](#), on page 99
- [Summary Address Tab](#), on page 47
- [Interface Tab](#), on page 66

Field Reference

Table 54: OSPF General Tab

Element	Description
	The General tab provides two identical sections; each is used to enable one OSPF process. The following options are available in each section.
Enable this OSPF Process	Check this box to enable an OSPF process. You cannot enable an OSPF process if you have RIP enabled on the security appliance. Deselect this option to remove the OSPF process.
OSPF Process ID	Enter a unique numeric identifier for the OSPF process. This process ID is used internally and does not need to match the OSPF process ID on any other OSPF devices. Valid values are from 1 to 65535.
Advanced button	Opens the OSPF Advanced Dialog Box , on page 82, in which you can configure additional process-related parameters, such as Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings.

OSPF Advanced Dialog Box

Use the OSPF Advanced dialog box to configure settings such as the Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings for an OSPF process.



Note Beginning with ASA version 9.2(1), certain OSPF settings have changed. If you configure a shared policy that uses settings specific to ASA 9.2(1)+, you will receive a validation error if that policy is assigned to a device whose version is earlier than 9.2(1). Likewise, if you configure a shared policy that uses settings that no longer apply to ASA 9.2(1)+, you will receive a validation error if that policy is assigned to an 9.2(1)+ device.

Navigation Path

You can access the OSPF Advanced dialog box from the [General Tab](#), on page 81.

Related Topics

- [Configuring OSPF](#), on page 80

Field Reference

Table 55: OSPF Advanced Dialog Box

Element	Description
OSPF Process	Displays the ID of the OSPF process you are configuring. You cannot change this value in this dialog box.
General Tab	
Router ID	To use a fixed router ID, select IP Address and then enter a router ID in IP address format in the Router ID field. To have the router ID automatically generated (the highest-level IP address on the security appliance is used as the router ID), select Automatic .
Ignore LSA MOSPF	Select this option to suppress transmission of syslog messages when the security appliance receives Type 6 (MOSPF) LSA packets.
RFC 1583 Compatible	Select this option to calculate summary route costs per RFC 1583. Deselect this option to calculate summary route costs per RFC 2328. To minimize the chance of routing loops, all OSPF devices in an OSPF routing domain should have RFC compatibility set identically. This option is selected by default.
Adjacency Changes	These options specify the syslog messages sent when adjacency changes occur. <ul style="list-style-type: none"> • Log Adjacency Changes – When selected, the security appliance sends a syslog message whenever an OSPF neighbor goes up or down. This option is selected by default. • Log Adjacency Changes Detail – When selected, the security appliance sends a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This option is not selected by default.

Element	Description
Administrative Route Distances	<p data-bbox="613 296 1403 323">Settings for the administrative route distances, according to the route type.</p> <ul data-bbox="651 342 1481 600" style="list-style-type: none"><li data-bbox="651 342 1481 405">• Inter Area – The administrative distance for all routes from one area to another. Valid values range from 1 to 255; the default value is 110.<li data-bbox="651 426 1481 489">• Intra Area – The administrative distance for all routes within an area. Valid values range from 1 to 255; the default value is 110.<li data-bbox="651 510 1481 600">• External – The administrative distance for all routes from other routing domains that are learned through redistribution. Valid values range from 1 to 255; the default value is 110.

Element	Description
Timers	

Element	Description
	<p>Settings used to configure LSA arrival, LSA pacing, and throttling for ASA 9.2(1)+ devices:</p> <ul style="list-style-type: none"> • LSA Arrival – The minimum delay in milliseconds that must pass between acceptance of the same LSA arriving from neighbors. The range is from 0 to 600,000 milliseconds. The default is 1000 milliseconds. • LSA Flood Pacing – The time in milliseconds at which LSAs in the flooding queue are paced in between updates. The configurable range is from 5 to 100 milliseconds. The default value is 33 milliseconds. • LSA Group Pacing – The interval at which LSAs are collected into a group and refreshed, checksummed, or aged. Valid values range from 10 to 1800; the default value is 240 seconds. • LSA Retransmission Pacing - The time in milliseconds at which LSAs in the retransmission queue are paced. The configurable range is from 5 to 200 milliseconds. The default value is 66 milliseconds. • LSA Throttle – The delay in milliseconds to generate the first occurrence of the LSA. Valid values range from 0 to 600000 milliseconds. When you enter a value in this field, the Min and Max fields are enabled: <ul style="list-style-type: none"> • Min – The minimum delay for originating the same LSA. Valid values range from 1 to 600000 milliseconds. • Max – The maximum delay for originating the same LSA. Valid values range from 1 to 600000 milliseconds. <p>Note For LSA throttling, the first occurrence value must be equal to or less than the minimum value and the minimum value must be equal to or less than the maximum value.</p> <ul style="list-style-type: none"> • SPF Throttle – The delay to receive a change to the SPF calculation. Valid values range from 1 to 600000 milliseconds. When you enter a value in this field, the Min and Max fields are enabled: <ul style="list-style-type: none"> • Min – The delay between the first and second SPF calculations. Valid values range from 1 to 600000 milliseconds. • Max – The maximum wait time for SPF calculations. Valid values range from 1 to 600000 milliseconds. <p>Note For SPF throttling, the first occurrence value must be equal to or less than the minimum value and the minimum value must be equal to or less than the maximum value.</p> <p>Settings used to configure LSA pacing and SPF calculation timers for device versions earlier than 9.2(1):</p> <ul style="list-style-type: none"> • SPF Delay – The time between receipt of a topology change and the start of shortest path first (SPF) calculations. Valid values range from 0 to 65535; the default value is 5 seconds.

Element	Description
	<ul style="list-style-type: none"> • SPF Hold – The hold time between consecutive SPF calculations. Valid values range from 1 to 65534; the default value is 10 seconds. • LSA Group Pacing – The interval at which LSAs are collected into a group and refreshed, checksummed, or aged. Valid values range from 10 to 1800; the default value is 240 seconds.
Default Information Originate	<p>Settings used by an ASBR to generate a default external route into an OSPF routing domain.</p> <ul style="list-style-type: none"> • Enable Default Information Originate – Check this box to enable generation of a default route into the OSPF routing domain; the following options become available: <ul style="list-style-type: none"> • Always advertise the default route – Check this box to always advertise the default route. • Metric Value – Enter the OSPF metric for the default route. Valid values range from 0 to 16777214; the default value is 1. • Metric Type – Choose the external link type associated with the default route advertised into the OSPF routing domain. The choices are 1 or 2, indicating a Type 1 or a Type 2 external route. The default value is 2. • Route Map – (Optional) Enter or Select a route map object to apply. The routing process generates the default route if the route map is satisfied. <p>Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.</p>
Non Stop Forwarding Tab	
<p>Note Non Stop Forwarding (NSF) is supported on ASA 9.3(1)+ devices in Spanned Cluster mode or Failover mode only.</p>	
Enable Cisco Non Stop Forwarding Capability	Enables configuration of Cisco nonstop forwarding (NSF) operations.
Enable Cisco Non Stop Forwarding Helper mode	<p>Enables Cisco nonstop forwarding (NSF) helper mode.</p> <p>When an ASA has NSF enabled, it is said to be NSF-capable and will operate in graceful restart mode--the OSPF router process performs nonstop forwarding recovery due to a Route Processor (RP) switchover. By default, the neighboring ASAs of the NSF-capable ASA will be NSF-aware and will operate in NSF helper mode. When the NSF-capable ASA is performing graceful restart, the helper ASAs assist in the nonstop forwarding recovery process.</p> <p>If you do not want the ASA to help the restarting neighbor with nonstop forwarding recovery, clear the Enable Cisco Non Stop Forwarding Helper mode option.</p>

Element	Description
Enable Cisco Non Stop Forwarding	Enables Cisco nonstop forwarding (NSF).
Cancel NSF restart when non-NSF-aware neighboring networking devices are detected (Enforce Global)	<p>If neighbors that are not NSF-aware are detected on a network interface during an NSF graceful restart, restart is aborted on that interface only and graceful restart will continue on other interfaces. To cancel restart for the entire OSPF process when neighbors that are not NSF-aware are detected during restart, select the Cancel NSF restart when non-NSF-aware neighboring networking devices are detected (Enforce Global) option.</p> <p>Note The NSF graceful restart will also be canceled for the entire process when a neighbor adjacency reset is detected on any interface or when an OSPF interface goes down.</p>
Enable IETF Non Stop Forwarding Capability	Enables configuration of Internet Engineering Task Force (IETF) NSF operations.
Enable IETF Non Stop Forwarding Helper mode	<p>Enables IETF nonstop forwarding (NSF) helper mode.</p> <p>When an ASA has NSF enabled, it is said to be NSF-capable and will operate in graceful restart mode--the OSPF router process performs nonstop forwarding recovery due to a Route Processor (RP) switchover. By default, the neighboring ASAs of the NSF-capable ASA will be NSF-aware and will operate in NSF helper mode. When the NSF-capable ASA is performing graceful restart, the helper ASAs assist in the nonstop forwarding recovery process.</p> <p>If you do not want the ASA to help the restarting neighbor with nonstop forwarding recovery, clear the Enable IETF Non Stop Forwarding Helper mode option.</p>
Enable Strict Link State advertisement checking	Enables strict link-state advertisement (LSA) checking for IETF NSF helper mode.
Enable IETF Non Stop Forwarding	Enables IETF nonstop forwarding (NSF).
Length of graceful restart interval	<p>(Optional) Specifies the length of the graceful restart interval, in seconds. The range is from 1 to 1800. The default is 120.</p> <p>Note For a restart interval below 30 seconds, graceful restart will be terminated.</p>

Area Tab

Use the Area tab on the OSPF page to configure OSPF areas and networks.

Navigation Path

You can access the Area tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#) , on page 80.

Related Topics

- [Add/Edit Area/Area Networks Dialog Box](#) , on page 89

- [Configuring OSPF](#) , on page 80
- [General Tab](#) , on page 81
- [Range Tab](#) , on page 91
- [Neighbors Tab](#) , on page 42
- [Redistribution Tab](#) , on page 44
- [Virtual Link Tab](#) , on page 96
- [Filtering Tab](#) , on page 99
- [Summary Address Tab](#) , on page 47
- [Interface Tab](#) , on page 66

Field Reference

Table 56: Area Tab

Element	Description
OSPF Process	The OSPF process the area applies to.
Area ID	The area ID.
Area Type	The area type (Normal, Stub, or NSSA).
Networks	The area networks.
Options	The options, if any, set for the area type.
Authentication	The type of authentication set for the area (None, Password, or MD5).
Cost	The default cost for the area.

Add/Edit Area/Area Networks Dialog Box

Use the Add/Edit Area/Area Networks dialog box to define area parameters, the networks contained by the area, and the OSPF process associated with the area.

Navigation Path

You can access the Add/Edit Area/Area Networks dialog box from the [Area Tab](#) , on page 88.

Related Topics

- [Configuring OSPF](#) , on page 80

Field Reference

Table 57: Add/Edit Area/Area Networks Dialog Box

Element	Description
OSPF Process	When adding a new area, choose the OSPF process ID for the OSPF process for which the area is being added. If there is only one OSPF process enabled on the security appliance, that process is selected by default. When editing an existing area, you cannot change the OSPF process ID.
Area ID	When adding a new area, enter the area ID. You can specify the area ID as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295. You cannot change the area ID when editing an existing area.
Area Type	
Normal	Choose this option to make the area a standard OSPF area. This option is selected by default when you first create an area.
Stub	Choosing this option makes the area a stub area. Stub areas do not have any routers or areas beyond it. Stub areas prevent AS External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you can prevent summary LSAs (Type 3 and 4) from being flooded into the area by deselecting the Summary check box.
Summary (allows sending LSAs into the stub area)	When the area being defined is a stub area, deselecting this check box prevents LSAs from being sent into the stub area. This check box is selected by default for stub areas.
NSSA	Choose this option to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create a NSSA, you can prevent summary LSAs from being flooded into the area by deselecting the Summary check box. You can also disable route redistribution by deselecting the Redistribute check box and enabling Default Information Originate.
Redistribute (imports routes to normal and NSSA areas)	Deselect this check box to prevent routes from being imported into the NSSA. This check box is selected by default.
Summary (allows sending LSAs into the NSSA area)	When the area being defined is a NSSA, deselecting this check box prevents LSAs from being sent into the stub area. This check box is selected by default for NSSAs.
Default Information Originate (generate a Type 7 default)	Select this check box to generate a Type 7 default into the NSSA. This check box is deselected by default.
Metric Value	Specifies the OSPF metric value for the default route. Valid values range from 0 to 16777214. The default value is 1.
Metric Type	The OSPF metric type for the default route. The choices are 1 (Type 1) or 2 (Type 2). The default value is 2.

Element	Description
Network	The IP address and network mask of the network or host to be added to the area. Use 0.0.0.0 with a netmask of 0.0.0.0 to create the default area. You can only use 0.0.0.0 in one area. Tip You can click Select to select the interfaces from a list of interface objects.
Authentication	Contains the settings for OSPF area authentication. <ul style="list-style-type: none"> • None—Choose this option to disable OSPF area authentication. This is the default setting. • Password—Choose this option to use a clear text password for area authentication. This option is not recommended where security is a concern. • MD5—Choose this option to use MD5 authentication.
Default Cost	Specify a default cost for the area. Valid values range from 0 to 65535 for ASA devices earlier than 9.2(1) and from 0 to 16777214 for ASA 9.2(1)+. The default value is 1.

Range Tab

Use the Range tab to summarize routes between areas.

Navigation Path

You can access the Range tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#), on page 80.

Related Topics

- [Add/Edit Area Range Network Dialog Box](#), on page 92

Field Reference

Table 58: Range Tab

Element	Description
Process ID	The ID of the OSPF process associated with the route summary.
Area ID	The ID of the area associated with the route summary.
Network	The summary IP address and network mask.
Advertise	Displays “true” if the route summaries are advertised when they match the address/mask pair or “false” if the route summaries are suppressed when they match the address/mask pair.

Add/Edit Area Range Network Dialog Box

Use the Add/Edit Area Range Network dialog box to add a new entry to the Route Summarization table or to change an existing entry.

Navigation Path

You can access the Add/Edit Area Range Network dialog box from the [Range Tab](#), on page 91.

Related Topics

- [Configuring OSPF](#), on page 80

Field Reference

Table 59: Add/Edit Area Range Network Dialog Box

Element	Description
OSPF Process	Select the OSPF process to which the route summary applies. You cannot change this value when editing an existing route summary entry.
Area	Select the area ID of the area to which the route summary applies. You cannot change this value when editing an existing route summary entry.
Network	The IP address and mask of the network for the routes being summarized. Tip You can click Select to select the networks from a list of network objects.
Advertise	Select this check box to set the address range status to “advertise”. This causes Type 3 summary LSAs to be generated. Deselect this check box to suppress the Type 3 summary LSA for the specified networks.

Neighbors Tab

Use the Neighbors tab to define static neighbors. You need to define a static neighbor for each point-to-point, non-broadcast interface. You also need to define a static route for each static neighbor in the Neighbors table.

Navigation Path

You can access the Neighbors tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#), on page 80.

Related Topics

- [Add/Edit Static Neighbor Dialog Box](#), on page 93

Field Reference

Table 60: Neighbors Tab

Element	Description
OSPF Process	The OSPF process associated with the static neighbor.
Neighbor	The IP address of the static neighbor.
Interface	The interface associated with the static neighbor.

Add/Edit Static Neighbor Dialog Box

Use the Add/Edit Static Neighbor dialog box to define a static neighbor or change information for an existing static neighbor. You must define a static neighbor for each point-to-point, non-broadcast interface.

Navigation Path

You can access the Add/Edit Static Neighbor dialog box from the [Neighbors Tab](#) , on page 92.

Related Topics

- [Configuring OSPF](#) , on page 80

Field Reference

Table 61: Add/Edit Static Neighbor Dialog Box

Element	Description
OSPF Process	The OSPF process associated with the static neighbor.
Neighbor	The IP address of the static neighbor. Tip You can click Select to select the neighbor from a list of host objects.
Interface	The interface associated with the static neighbor. Tip You can click Select to select the interface from a list of interface objects.

Redistribution Tab

Use the Redistribution tab to define the rules for redistributing routes from one routing domain to another.

Navigation Path

You can access the Redistribution tab from the OSPF Page. For more information about the OSPF page, see [Configuring OSPF](#) , on page 80.

Related Topics

- [Redistribution Dialog Box](#) , on page 94

Field Reference**Table 62: Redistribution Tab**

Element	Description
OSPF Process	The OSPF process associated with the route redistribution entry.
Route Type	The source protocol the routes are being redistributed from. Valid entries are the following: <ul style="list-style-type: none"> • BGP—Redistribute routes from the BGP routing process. • Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the AS. • EIGRP—Redistributes routes from the EIGRP routing process. Choose the autonomous system number of the EIGRP routing process from the list. • OSPF—Redistributes routes from another OSPF routing process. • RIP—Redistributes routes from the RIP routing process. • Static—Redistributes static routes to the OSPF routing process.
Match	The conditions used for redistributing routes from one routing protocol to another. These options are not available when redistributing static, connected, RIP, BGP, or EIGRP routes.
Subnets	Displays “true” if subnetted routes are redistributed. Does not display anything if only routes that are not subnetted are redistributed.
Metric Value	The metric that is used for the route. This column is blank for redistribution entries if the default metric is used.
Metric Type	Displays “1” if the metric is a Type 1 external route, “2” if the metric is Type 2 external route.
Tag Value	A 32-bit decimal value attached to each external route. This value is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.
Route Map	The name of the route map object to apply to the redistribution entry.

Redistribution Dialog Box

Use the Redistribution dialog box to add a redistribution rule or to edit an existing redistribution rule in the Redistribution table.

Navigation Path

You can access the Redistribution dialog box from the [Redistribution Tab](#) , on page 93.

Related Topics

- [Configuring OSPF](#) , on page 80

Field Reference**Table 63: OSPF Redistribution Settings Dialog Box**

Element	Description
OSPF Process	Select the OSPF process associated with the route redistribution entry.
Route Type	Select the source protocol from which the routes are being redistributed. You can choose one of the following options: <ul style="list-style-type: none"> • BGP—Redistribute routes from the BGP routing process. • Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the AS. • EIGRP—Redistributes routes from the EIGRP routing process. Choose the autonomous system number of the EIGRP routing process from the list. • OSPF—Redistributes routes from another OSPF routing process. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, BGP, or EIGRP routes. • RIP—Redistributes routes from the RIP routing process. • Static—Redistributes static routes to the OSPF routing process.
Routing Process ID	The autonomous system (AS) number for the BGP or EIGRP routing process.
Match	If you have chosen OSPF as the Route Type, choose the conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions: <ul style="list-style-type: none"> • Internal—The route is internal to a specific AS. • External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes. • External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes. • NSSA External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes. • NSSA External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.

Element	Description
Metric Value	The metric value for the routes being redistributed. Valid values range from 1 to 16777214. When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.
Metric Type	Select “1” if the metric is a Type 1 external route, “2” if the metric is a Type 2 external route.
Tag Value	The tag value is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.
Use Subnets	When selected, redistribution of subnetted routes is enabled. Deselect this check box to cause only routes that are not subnetted to be redistributed.
Route Map	Enter or Select a route map object to apply to the redistribution entry. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.

Virtual Link Tab

Use the Virtual Link tab to create virtual links. If you add an area to an OSPF network, and it is not possible to connect the area directly to the backbone area, you need to create a virtual link. A virtual link connects two OSPF devices that have a common area, called the transit area. One of the OSPF devices must be connected to the backbone area.

Navigation Path

You can access the Virtual Link tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#) , on page 80.

Related Topics

- [Add/Edit OSPF Virtual Link Configuration Dialog Box](#) , on page 97

Field Reference

Table 64: Virtual Link Tab

Element	Description
OSPF Process	The OSPF process associated with the virtual link.
Area ID	The ID of the transit area.
Peer Router	The IP address of the virtual link neighbor.

Element	Description
Authentication	Displays the type of authentication used by the virtual link: <ul style="list-style-type: none"> • None—No authentication is used. • Password—Clear text password authentication is used. • MD5—MD5 authentication is used. • Key Chain—Key chain authentication is enabled.

Add/Edit OSPF Virtual Link Configuration Dialog Box

Use the Add/Edit OSPF Virtual Link Configuration dialog box to define virtual links or change the properties of existing virtual links.

Navigation Path

You can access the Add/Edit OSPF Virtual Link Configuration dialog box from the [Virtual Link Tab](#) , on [page 96](#).

Related Topics

- [Add/Edit OSPF Virtual Link MD5 Configuration Dialog Box](#) , on [page 99](#)
- [Configuring OSPF](#) , on [page 80](#)

Field Reference

Table 65: Add/Edit OSPF Virtual Link Configuration Dialog Box

Element	Description
OSPF Process	Select the OSPF process associated with the virtual link.
Area ID	Select the area shared by the neighbor OSPF devices. The selected area cannot be an NSSA or a stub area.
Peer Router	Enter the IP address of the virtual link neighbor.
Hello Interval	The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds for ASA devices earlier than 9.2(1) and from 1 to 8192 seconds for ASA 9.2(1)+. The default value is 10 seconds.

Element	Description
Retransmit Interval	The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds for ASA devices earlier than 9.2(1) and from 1 to 8192 seconds for ASA 9.2(1)+. The default value is 5 seconds.
Transmit Delay	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds for ASA devices earlier than 9.2(1) and from 1 to 8192 seconds for ASA 9.2(1)+. The default value is 1 second.
Dead Interval	The interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535 seconds for ASA devices earlier than 9.2(1) and from 1 to 8192 seconds for ASA 9.2(1)+. The default value of this field is four times the interval set by the Hello Interval field.
Authentication	<p>Contains the OSPF authentication options.</p> <ul style="list-style-type: none"> • None—Choose this option to disable OSPF authentication. • Area—Choose this option to use the authentication type specified for the area. See Add/Edit Area/Area Networks Dialog Box, on page 89 for information about configuring area authentication. Area authentication is disabled by default. Therefore, unless you have previously specified an area authentication type, interfaces set to area authentication have authentication disabled until you configure this setting. • Password—Choose this option to use clear text password authentication. This is not recommended where security is a concern. • MD5—Choose this option to use MD5 authentication (recommended). • Key Chain—Choose this option to use key chain authentication.
Key Chain	<p>This field appears when Key Chain authentication is enabled. Click Select and choose the configured key chain. To know the configuration steps, refer Configuring Key Chain, on page 108.</p> <p>Note Use the same authentication type and key ID for the peers to establish a successful adjacency.</p>
Authentication Password	<p>Contains the settings for entering the password when password authentication is enabled.</p> <ul style="list-style-type: none"> • Password—Enter a text string of up to 8 characters. • Confirm—Re-enter the password.

Element	Description
MD5 IDs and Keys	<p>Contains the settings for entering the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID.</p> <ul style="list-style-type: none"> • MD5 Key ID and MD5 Key Table <ul style="list-style-type: none"> • MD5 Key ID—A numerical key identifier. Valid values range from 1 to 255. • MD5 Key—An alphanumeric character string of up to 16 bytes.

Add/Edit OSPF Virtual Link MD5 Configuration Dialog Box

Use the Add/Edit OSPF Virtual Link MD5 Configuration dialog box to define MD5 keys for authentication of virtual links.

Navigation Path

You can access the Add/Edit OSPF Virtual Link MD5 Configuration dialog box from the [Add/Edit OSPF Virtual Link Configuration Dialog Box](#) , on page 97.

Related Topics

- [Add/Edit OSPF Virtual Link Configuration Dialog Box](#) , on page 97
- [Virtual Link Tab](#) , on page 96
- [Configuring OSPF](#) , on page 80

Field Reference

Table 66: Add/Edit OSPF Virtual Link MD5 Configuration Dialog Box

Element	Description
MD5 Key ID	A numerical key identifier. Valid values range from 1 to 255.
MD5 Key	An alphanumeric character string of up to 16 bytes.
Confirm	Re-enter the MD5 key.

Filtering Tab

Use the Filtering tab to configure the ABR Type 3 LSA filters for each OSPF process. ABR Type 3 LSA filters allow only specified prefixes to be sent from one area to another area and restricts all other prefixes. This type of area filtering can be applied out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF areas at the same time.

Benefits

OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Restrictions

Only type-3 LSAs that originate from an ABR are filtered.

Navigation Path

You can access the Filtering tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#) , on page 80.

Related Topics

- [Add/Edit Filtering Dialog Box](#) , on page 100

Field Reference

Table 67: Filtering Tab

Element	Description
OSPF Process	The OSPF process associated with the filter entry.
Area ID	The ID of the area associated with the filter entry.
Prefix List Name	The name of the prefix list.
Filtered Network	The IP address and mask of the network being filtered.
Traffic Direction	Displays “Inbound” if the filter entry applies to LSAs coming in to an OSPF area or “Outbound” if it applies to LSAs going out of an OSPF area.
Sequence #	The sequence number for the filter entry. When multiple filters apply to an LSA, the filter with the lowest sequence number is used.
Action	Displays “Permit” if LSAs matching the filter are allowed or “Deny” if LSAs matching the filter are denied.
Lower Range	The minimum prefix length to be matched.
Upper Range	The maximum prefix length to be matched.

Add/Edit Filtering Dialog Box

Use the Add/Edit Filtering dialog box to add new filters to the Filter table or to modify an existing filter.

Navigation Path

You can access the Add/Edit Filtering dialog box from the [Filtering Tab](#) , on page 99.

Related Topics

- [Configuring OSPF](#) , on page 80

Field Reference

Table 68: Add/Edit Filtering Dialog Box

Element	Description
OSPF Process	Select the OSPF process associated with the filter entry.
Area ID	Select the ID of the area associated with the filter entry.
Prefix List Name	Enter or Select the appropriate prefix list object. Tip Click Select to open the Prefix List Object Selector from which you can select a prefix list object. You can also create new objects from the object Prefix List Object selector. For more information, see Add or Edit Prefix List Object Dialog Box , on page 159.
Filtered Network	Enter the IP address and mask of the network being filtered.
Traffic Direction	Select the traffic direction to filter. Choose “Inbound” to filter LSAs coming into an OSPF area or “Outbound” to filter LSAs going out of an OSPF area.
Sequence Number	Enter a sequence number for the filter. Valid values range from 1 to 4294967294. When multiple filters apply to an LSA, the filter with the lowest sequence number is used.
Action	Select “Permit” to allow the LSA traffic or “Deny” to block the LSA traffic.
Lower Range	Specify the minimum prefix length to be matched. The value of this setting must be greater than the length of the network mask entered in the Filtered Network field and less than or equal to the value, if present, entered in the Upper Range field.
Upper Range	Enter the maximum prefix length to be matched. The value of this setting must be greater than or equal to the value, if present, entered in the Lower Range field, or, if the Lower Range field is left blank, greater than the length of the network mask length entered in the Filtered Network field.

Filter Rule Tab

Use the Filter Rule tab to configure rules to filter networks received or transmitted in Open Shortest Path First (OSPF) updates.



Note Filter rules are supported on ASA 9.2(1)+ only.

Navigation Path

You can access the Filter Rule tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#) , on page 80.

Related Topics

- [Add/Edit Filter Rule Dialog Box](#) , on page 102

Field Reference*Table 69: Filter Rule Tab*

Element	Description
Process ID	The OSPF process associated with the filter rule.
ACL	Standard IP access list name. The list defines which networks are to be received and which are to be suppressed in routing updates.
Direction	The direction for the filter rule: <ul style="list-style-type: none"> • in—The rule filters default route information from incoming routing updates. • out—The rule filters default route information from outgoing routing updates.
Interface	(Optional) The interface to which the filter rule applies.
Routing Process	The routing process: None, BGP, Connected, EIGRP, OSPF, RIP, or Static.
Routing Process ID	The identifier for the routing process.

Add/Edit Filter Rule Dialog Box

Use the Add/Edit Filter Rule dialog box to add new filter rules to the Filter Rules table or to modify an existing filter rule.



Note Filter rules are supported on ASA 9.2(1)+ only.

Navigation Path

You can access the Add/Edit Filter Rule dialog box from the [Filter Rule Tab](#) , on page 101.

Related Topics

- [Configuring OSPF](#) , on page 80

Field Reference

Table 70: Add/Edit Filter Rule Dialog Box

Element	Description
OSPF Process	Select the OSPF process associated with the filter rule.
ACL	Select an Access Control List that defines which networks are to be received and which are to be suppressed in routing updates.
Direction	Specify the direction for the filter rule: <ul style="list-style-type: none"> • in—The rule filters default route information from incoming routing updates. • out—The rule filters default route information from outgoing routing updates.
Interface	(Optional) Specify the interface on which to apply the routing updates. Specifying an interface causes the access list to be applied only to routing updates received on that interface.
Routing Process	Select the routing process for which you want to filter: None, BGP, Connected, EIGRP, OSPF, RIP, or Static.
Routing Process ID	Enter the identifier for the routing process. Applies to BGP, EIGRP, and OSPF routing protocols.

Summary Address Tab

Use the Summary Address tab to configure summary addresses for each OSPF routing process.

Routes learned from other routing protocols can be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. Summary routes help reduce the size of the routing table.

Using summary routes for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address. Only routes from other routing protocols that are being redistributed into OSPF can be summarized.

Navigation Path

You can access the Summary Address tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#), on page 80.

Related Topics

- [Add/Edit Summary Address Dialog Box](#), on page 104

Field Reference

Table 71: Summary Address Tab

Element	Description
Process ID	The OSPF process associated with the summary address.
Network	The IP address and network mask of the summary address.
Tag	A 32-bit decimal value attached to each external route. This value is not used by OSPF itself. It may be used to communicate information between ASBRs.
Advertise	Displays “true” if the summary routes are advertised. Displays “false” if the summary route is not advertised.

Add/Edit Summary Address Dialog Box

Use the Add/Edit Summary Address dialog box to add new entries or to modify existing entries in the Summary Address table.

Navigation Path

You can access the Add/Edit Summary Address dialog box from the [Summary Address Tab](#), on page 103.

Related Topics

- [Configuring OSPF](#), on page 80

Field Reference

Table 72: Add/Edit Summary Address Dialog Box

Element	Description
OSPF Process	Choose the OSPF process associated with the summary address. You cannot change this information when editing an existing entry.
Network	The IP address and network mask of the summary address.
Tag	The tag value is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.
Advertise	When selected, summary routes are advertised. Deselect this check box to suppress routes that fall under the summary address. By default, this check box is selected.

Interface Tab

Use the Interface tab to configure interface-specific OSPF authentication routing properties.

Navigation Path

You can access the Interface tab from the OSPF page. For more information about the OSPF page, see [Configuring OSPF](#), on page 80.

Related Topics

- [Add/Edit Interface Dialog Box](#), on page 106

Field Reference

Table 73: Interface Tab

Element	Description
Interface	The name of the interface to which the configuration applies.
Authentication	The type of OSPF authentication enabled on the interface. The authentication type can be one of the following values: <ul style="list-style-type: none"> • None—OSPF authentication is disabled. • Password—Clear text password authentication is enabled. • MD5—MD5 authentication is enabled. • Area—The authentication type specified for the area is enabled on the interface. Area authentication is the default value for interfaces. However, area authentication is disabled by default. So, unless you previously specified an area authentication type, interfaces showing Area authentication have authentication disabled. • Key Chain—Key chain authentication is enabled.
Point-to-Point	Displays “true” if the interface is set to non-broadcast (point-to-point). Displays “false” if the interface is set to broadcast.
Cost	The cost of sending a packet through the interface.
Priority	The OSPF priority assigned to the interface.
MTU Ignore	Displays “false” if MTU mismatch detection is enabled. Displays “true” if the MTU mismatch detection is disabled.
Database Filter	Displays “true” if outgoing LSAs are filtered during synchronization and flooding. Displays “false” if filtering is not enabled.
Hello Interval	The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds. The default value is 10 seconds.

Element	Description
Transmit Delay	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds. The default value is 1 second.
Retransmit Interval	The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Dead Interval	The interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535. The default value of this setting is four times the interval set by the Hello Interval field.
Hello Multiplier (ASA 9.2(1)+ only)	The number of hello packets to be sent per second. Valid values are between 3 and 20.

Add/Edit Interface Dialog Box

Use the Add/Edit Interface dialog box to add OSPF authentication routing properties for an interface or to change an existing entry.



Note Beginning with ASA version 9.2(1), the upper limit for acceptable entries for Hello Interval, Transmit Delay, Retransmit Interval, and Dead Interval has been reduced from 65535 seconds to 8192 seconds. If you configure a shared policy that uses a value over 8192, you will receive a validation error if that policy is assigned to an 9.2(1)+ device.

Navigation Path

You can access the Add/Edit Interface dialog box from the [Interface Tab](#) , on page 66.

Related Topics

- [Configuring OSPF](#) , on page 80

Field Reference

Table 74: Add/Edit Interface Dialog Box

Element	Description
Interface	The name of the interface to which the configuration applies.

Element	Description
Authentication	<p>The type of OSPF authentication enabled on the interface. The authentication type can be one of the following values:</p> <ul style="list-style-type: none"> • No Authentication—OSPF authentication is disabled. • Area Authentication—The authentication type specified for the area is enabled on the interface. Area authentication is the default value for interfaces. However, area authentication is disabled by default. So, unless you previously specified an area authentication type, interfaces showing Area authentication have authentication disabled. • Password Authentication—Clear text password authentication is enabled. • MD5 Authentication—MD5 authentication is enabled. • Key Chain—Key chain authentication is enabled.
Key Chain	<p>Click Select and choose the configured key chain. To know the configuration steps, refer Configuring Key Chain, on page 108.</p> <p>Note Use the same authentication type and key ID for the peers to establish a successful adjacency.</p>
Authentication Password	<p>Contains the settings for entering the password when password authentication is enabled.</p> <ul style="list-style-type: none"> • Enter Password—Enter a text string of up to 8 characters. • Confirm—Re-enter the password.
MD5 Key IDs and Keys	<p>Contains the settings for entering the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID.</p> <ul style="list-style-type: none"> • Key ID—Enter a numerical key identifier. Valid values range from 1 to 255. • Key—An alphanumeric character string of up to 16 bytes. • Confirm—Re-enter the MD5 key. <p>Enter the above values, then click >> to add the key information to the Keys table. Select a key entry and then click << to remove it from the Keys table.</p>
Cost	The cost of sending a packet through the interface.
Priority	The OSPF priority assigned to the interface.
MTU Ignore	When selected, MTU mismatch detection is disabled. Clear this check box to enable MTU mismatch detection.
Database Filter All Out	When selected, outgoing LSAs are filtered during synchronization and flooding. Deselect this check box to disable filtering.

Element	Description
Hello Interval (sec)	<p>The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface.</p> <p>For ASA 9.2(1)+ devices, valid values range from 1 to 8192 seconds. For all other devices, valid values range from 1 to 65535 seconds. The default value is 10 seconds.</p>
Transmit Delay (sec)	<p>The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links.</p> <p>For ASA 9.2(1)+ devices, valid values range from 1 to 8192 seconds. For all other devices, valid values range from 1 to 65535 seconds. The default value is 1 second.</p>
Retransmit Interval (sec)	<p>The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links.</p> <p>For ASA 9.2(1)+ devices, valid values range from 1 to 8192 seconds. For all other devices, valid values range from 1 to 65535 seconds. The default value is 5 seconds.</p>
Dead Interval (sec)	<p>The interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down.</p> <p>For ASA 9.2(1)+ devices, valid values range from 1 to 8192 seconds. For all other devices, valid values range from 1 to 65535 seconds. The default value of this setting is four times the interval set by the Hello Interval field.</p>
Hello Multiplier (Hello/Sec) (ASA 9.2(1)+ only)	<p>The number of hello packets to be sent per second. Valid values are between 3 and 20.</p> <p>Note If you specify a Hello Multiplier, the Hello Interval and Dead Interval values will be ignored. If you entered a value for Hello Interval or Dead Interval, you will be asked to confirm that you want to use the Hello Multiplier instead of the Hello Interval and Dead Interval settings.</p>
Point-to-Point	<p>Displays “true” if the interface is set to non-broadcast (point-to-point). Displays “false” if the interface is set to broadcast.</p>

Configuring Key Chain

To enhanced data security and protection on the networking devices, the devices are configured with rotating keys for authenticating IGP peers that have a duration of 180 days or less. The rotating keys prevent any malicious user from guessing the keys used for routing protocol authentication and thereby protecting the network from advertising incorrect routes and redirecting traffic. Changing the keys frequently reduces the

risk of them eventually being guessed. When configuring authentication for routing protocols that provide key chains, configure the keys in a key chain to have overlapping lifetimes. This configuration helps to prevent loss of key-secured communication due to absence of an active key. If the key lifetime expires and no active keys are found, OSPF uses the last valid key to maintain the adjacency with peers.

The two limitations of key chain configuration in Cisco Security Manager are:

- The configured Key ID will be displayed in unencrypted format in the [OOB \(Out of Band\) Changes Dialog Box](#).
- The option to copy provision is not available for key chains.

Related Topics

- [Lifetime of a Key](#), on page 109
- [Add/Edit Key Chain](#), on page 110

Lifetime of a Key

To maintain stable communications, each device stores key chain authentication keys and uses more than one key for a feature at the same time. Based on the send and accept lifetimes of a key, keychain management provides a secured mechanism to handle key rollover. The device uses the lifetimes of keys to determine which keys in a key chain are active.

Each key in a key chain has two lifetimes:

- Accept lifetime—The time interval within which the device accepts the key during key exchange with another device.
- Send lifetime—The time interval within which the device sends the key during key exchange with another device.

During a key send lifetime, the device sends routing update packets with the key. The device does not accept communication from other devices when the key sent is not within the accept lifetime of the key on the device.

If lifetimes are not configured then it is equivalent to configuring MD5 authentication key without timelines.

Key Selection

- When key chain has more than one valid key, OSPF selects the key that has the maximum life time.
- Key having an infinite lifetime is preferred.
- If keys have the same lifetime, then key with the higher key ID is preferred.

Related Topics

- [Configuring Key Chain](#), on page 108
- [Add/Edit Key Chain](#), on page 110

Add/Edit Key Chain

Use the Add/Edit KeyChain dialog box to add new entries or to modify existing entries in the KeyChain table.

Navigation Path

- You can access the Key Chain page tab from the Interface tab of OSPF page. For more information about the Interface tab, see [Interface Tab](#), on page 66.
- You can directly access the Add Key Chain page from **Manage > Policy Objects > Key Chain**.

Step 1 Create a key chain policy object that includes the key chains for authentication.

- Select **Manage > Policy Objects** to open the Policy Object Manager window (see [Policy Object Manager](#)).
- Select **Key Chain** from the table of contents.
- Right-click and choose **New Object**.
- In the Add Key Chain dialog box, enter a name for the object, for example, Chain 1.
- Click the **Add button** to add the key chain entry to the Key Chain list.

Step 2 Enter the relevant values in the **Add Key Chain Entry dialog box**:

Field Reference

Table 75: Add Key Chain Entry page

Element	Description
Algorithm	MD5 is the default cryptographic algorithm used for authentication.
Key ID	Enter a value between 0 and 255. Note The Key ID does not get displayed in encrypted format in the OOB (Out of Band) Changes Dialog Box .
Authentication Type	Select the relevant option: <ul style="list-style-type: none"> • Clear Text—To have the authentication key in text format. • Encryption —To have the authentication key in encrypted format.
Key String	Enter the key string.
Confirm Key String	Re-enter the same key string.
Accept Lifetime Settings— Provide the interval within which the device accepts the key during key exchange with another device.	
Timezone	Select either UTC or Local.
Start Date/Time	Provide the start date and the time in hh:mm:ss format.

Element	Description
End Time Type	Select the relevant option: <ul style="list-style-type: none"> • Date Time—The absolute time that the lifetime ends. • Duration—The number of seconds after the start time that the lifetime ends. • Infinite—Infinite lifetime (no end-time).
End Date	Provide the absolute date and the time. This option is not available if you choose Duration or Infinite as the End Time Type.
Duration	Provide the value in seconds after the start time that the lifetime ends. The permissible range is 1 and 2147483646. This option is not available if you choose Date Time or Infinite as the End Time Type.
Send Lifetime Settings—The time interval within which the device sends the key during key exchange with another device.	
Timezone	Select either UTC or Local.
Start Date/Time	Provide the start date and the time in hh:mm:ss format.
End Time Type	Select the relevant option: <ul style="list-style-type: none"> • Date Time—The absolute time that the lifetime ends. • Duration—The number of seconds after the start time that the lifetime ends. • Infinite—Infinite lifetime (no end-time).
End Date	Provide the absolute date and the time. This option is not available if you choose Duration or Infinite as the End Time Type.
Duration	Provide the value in seconds after the start time that the lifetime ends. The permissible range is 1 and 2147483646. This option is not available if you choose Date Time or Infinite as the End Time Type.

Step 3 Click Ok. Remember to submit your changes to the databases.

What to do next

Related Topics

- [Configuring Key Chain](#) , on page 108
- [Lifetime of a Key](#), on page 109

Configuring OSPFv3

The OSPFv3 page provides two tabbed panels for configuring OSPF (Open Shortest Path First) version 3 routing on a firewall device.

Navigation Path

- (Device view) Select **Platform > Routing > OSPFv3** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > OSPFv3** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

This is the basic procedure for configuring an OSPFv3 process and assigning it to an interface on the OSPFv3 page:

1. On the [Process Tab](#), on page 114:
 - Specify which of the two processes you are configuring by choosing **Process 1** or **Process 2** from the OSPFv3 Process drop-down list.
 - Check **Enable OSPFv3 Process**.
 - Assign a **Process ID**; any positive integer between 1 and 65535.
 - Use the following features as needed to define the process:
 - **Advanced** button, opening the [OSPFv3 Advanced Properties Dialog Box](#), on page 115.
 - **Area Tab (OSPFv3)**, on page 119, for managing area, range, and virtual-link definitions, by means of the [Add/Edit Area Dialog Box \(OSPFv3\)](#), on page 119, [Add/Edit Range Dialog Box \(OSPFv3\)](#), on page 121, and [Add/Edit Virtual Link Dialog Box \(OSPFv3\)](#), on page 122.
 - **Redistribution** panel, for managing route redistribution definitions by means of the [Add/Edit Redistribution Dialog Box \(OSPFv3\)](#), on page 123.
 - **Summary Prefix** panel, for managing summary-prefix definitions by means of the [Add/Edit Summary Prefix Dialog Box \(OSPFv3\)](#), on page 124.
2. On the [OSPFv3 Interface Tab](#), on page 125:
 - a. Use the Interface and Neighbor panels to assign the process to a specific interface, using the [Add/Edit Interface Dialog Box \(OSPFv3\)](#), on page 126 and the [Add/Edit Neighbor Dialog Box \(OSPFv3\)](#), on page 129.

Related Topics

- [About OSPFv3](#), on page 112

About OSPFv3

Open Shortest Path First (OSPF) is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. Version 3 is basically OSPFv2 enhanced for IPv6. It is similar to OSPFv2 (see

[About OSPF , on page 81](#)), but it is not backward compatible. To use OSPF to route both IPv4 and IPv6v packets, it will be necessary to run both OSPFv2 and OSPFv3 concurrently. They co-exist with each other, but do not interact.



Note OSPFv3 is supported on ASA 9.0+ devices operating in single-context, routed mode only. That is, multiple contexts and transparent mode are not supported.

Think of a link as being an interface on a networking device. A link-state protocol makes its routing decisions based on the states of the links that connect source and destination devices. The state of a link is a description of that interface and its relationship to its neighboring networking devices. This interface information includes the IPv6 prefix/length of the interface, the type of network it is connected to, the devices connected to that network, and so on. This information is propagated in various type of link-state advertisements (LSAs). Because only LSAs are exchanged, rather than entire routing tables, OSPF networks converge more quickly than RIP networks.

The ASA can run two processes of the OSPFv3 protocol simultaneously on different sets of interfaces. You might want to run two processes if you have interfaces that use the same IP addresses (NAT allows these interfaces to co-exist, but OSPFv3 does not allow overlapping addresses). Or you might want to run one process on the inside interface and another on the outside, redistributing a subset of routes between the two processes. Similarly, you might need to segregate private addresses from public addresses.

You can redistribute routes into an OSPFv3 routing process from another OSPFv3 routing process, a RIP routing process, or from static and connected routes configured on OSPFv3-enabled interfaces.

If NAT is employed but OSPFv3 is only running in public areas, routes to public networks can be redistributed inside the private network, either as default or type 5 AS External LSAs. However, you need to configure static routes for the private networks protected by the security appliance. Also, you should not mix public and private networks on the same security appliance interface.

Differences Between OSPFv2 and OSPFv3

The additional features provided by OSPFv3 over OSPFv2 include the following:

- Use of the IPv6 link-local address for neighbor discovery and other features.
- LSAs expressed as prefix and prefix length.
- Addition of two LSA types.
- Handling of unknown LSA types.
- Protocol processing per link.
- Removal of addressing semantics.
- Addition of flooding scope.
- Support for multiple instances per link.
- Authentication support using the IPsec ESP standard for OSPFv3 routing protocol traffic, as specified by RFC-4552.

Configuration Restrictions

The following are ASA OSPFv3 configuration restrictions:

- To enable OSPFv3 on a specific interface, IPv6 should be enabled on the interface and it must be named.
- Only one OSPFv3 process, with one area and one instance, can be assigned to an interface.
- The Interface neighbor entries take effect only when the OSPFv3 is enabled, and network type should be point-to-point on the specified interface.
- Interface neighbor address must be a link-local address.
- Range value in area Range table should be unique across the area.
- If the area is set to NSSA or stub, the same area cannot be set for virtual-link.
- OSPFv3 redistribution not applicable on the same OSPFv3 process.
- If used in an ASA cluster, OSPFv3 encryption should be disabled.
- The Layer 3 cluster pool is not shared between OSPFv3 and the interface.

Related Topics

- [Configuring OSPFv3](#) , on page 112
- [Process Tab](#) , on page 114
- [OSPFv3 Interface Tab](#) , on page 125

Process Tab

Use the Process tab on the OSPFv3 page to enable and configure up to two OSPFv3 routing processes. Each OSPF process has its own associated areas and networks. For each, at minimum, create an area for OSPFv3, enable an interface for OSPFv3, then redistribute the route into the targeted OSPFv3 routing processes. Note that only single-context mode is supported.

Navigation Path

The Process tab is on the OSPFv3 page.

- (Device view) Select **Platform > Routing > OSPFv3** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > OSPFv3** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- [Configuring OSPFv3](#) , on page 112
- [About OSPFv3](#) , on page 112
- [Area Tab \(OSPFv3\)](#) , on page 119
- [OSPFv3 Interface Tab](#) , on page 125

Field Reference

Table 76: Process Tab

Element	Description
OSPFv3 Process	Identify which OSPFv3 process you are configuring: choose Process 1 or Process 2 . You can enable one or both.
Enable OSPFv3 Process	Check this box to enable the chosen OSPFv3 process. Deselect this option to disable the OSPFv3 process; the process configuration information is retained should you wish to re-enable it later.
Process ID	Enter a unique numeric identifier for this process. The ID can be any positive integer between 1 and 65535. This process ID is used internally and does not need to match the OSPFv3 process ID on any other OSPFv3 devices.
Advanced	Opens the OSPFv3 Advanced Properties Dialog Box , on page 115, in which you can configure additional process-related parameters, such Router ID, Adjacency Changes, Administrative Route Distances, Timers, Default Information Originate, and Passive Interface settings.
Area	Use the tabs and tables in this panel to manage area, range and virtual-link definitions. See Area Tab (OSPFv3) , on page 119 for more about these definitions.
Redistribution	Use this panel to manage redistribution definitions. See Add/Edit Redistribution Dialog Box (OSPFv3) , on page 123 for more about these definitions.
Summary Prefix	Use this panel to manage summary prefix definitions. See Add/Edit Summary Prefix Dialog Box (OSPFv3) , on page 124 for more about these definitions.

OSPFv3 Advanced Properties Dialog Box

Use the OSPF Advanced dialog box to configure settings such as the Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings for an OSPF process.

Navigation Path

You can access the OSPF Advanced dialog box from the [Process Tab](#), on page 114.

Related Topics

- [Configuring OSPFv3](#), on page 112
- [About OSPFv3](#), on page 112

Field Reference

Table 77: OSPF Advanced Dialog Box

Element	Description
OSPF Process	This read-only field displays the ID of the OSPF process you are configuring.
Router ID	<p>On a single device, choose Automatic or IP Address. (An address field appears when you choose IP Address.)</p> <p>If you choose Automatic, the highest-level IP address on the security appliance is used as the router ID. To use a fixed router ID, choose IP Address and enter an IPv4 address in the Router ID field.</p> <p>On a device cluster, choose Automatic or Cluster Pool. (An IPv4 Pool object ID field appears when you choose Cluster Pool.)</p> <p>If you choose Cluster Pool, enter or Select the name of the IPv4 Pool object that is to supply the Router ID address. For more information, see Add or Edit IPv4 Pool Dialog Box.</p>
Ignore LSA MOSPF	Select this option to suppress transmission of syslog messages when the security appliance receives Type 6 (MOSPF) LSA packets.
Adjacency Changes	<p>These options specify the syslog messages sent when adjacency changes occur:</p> <ul style="list-style-type: none"> • Log Adjacency Changes – When selected, the security appliance sends a syslog message whenever an OSPF neighbor goes up or down. Checking this box enables the Include Details option. • Include Details – When selected, the security appliance sends a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This option is available only when Log Adjacency Changes is checked.
Administrative Route Distances	<p>Settings for the administrative route distances, according to the route type.</p> <ul style="list-style-type: none"> • Inter Area – The administrative distance for all routes from one area to another. Valid values range from 1 to 254; the default value is 110. • Intra Area – The administrative distance for all routes within an area. Valid values range from 1 to 254; the default value is 110. • External – The administrative distance for all routes from other routing domains that are learned through redistribution. Valid values range from 1 to 254; the default value is 110.

Element	Description
Timers (in milliseconds)	<p>LSA and SPF throttling provide a dynamic mechanism to slow LSA updates in OSPFv3 during times of network instability, and allow faster OSPFv3 convergence by providing LSA rate limiting. The settings used to configure LSA pacing and SPF calculation timers are:</p> <ul style="list-style-type: none"> • LSA Arrival – The minimum delay between acceptance of the same LSA arriving from neighbors. Valid values range from 0 to 600000 milliseconds. The default is 1000. • LSA Flood Pacing – The amount of time LSAs in the flooding queue are paced in between updates. Valid values range from 5 to 100 milliseconds. The default value is 33. • LSA Group Pacing – The interval at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800; the default value is 240 milliseconds. • LSA Retransmission Pacing – The length of time at which LSAs in the retransmission queue are paced. Valid values range from 5 to 200 milliseconds. The default value is 66. • LSA Throttle – The delay in milliseconds to generate the first occurrence of the LSA. Valid values range from 0 to 600000 milliseconds. When you enter a value in this field, the min and max fields are enabled: <ul style="list-style-type: none"> • min – The minimum delay for originating the same LSA. Valid values range from 1 to 600000 milliseconds. • max – The maximum delay for originating the same LSA. Valid values range from 1 to 600000 milliseconds. • SPF Throttle – The delay to receive a change to the SPF calculation. Valid values range from 1 to 600000 milliseconds. When you enter a value in this field, the min and max fields are enabled: <ul style="list-style-type: none"> • min – The delay between the first and second SPF calculations. Valid values range from 1 to 600000 milliseconds. • max – The maximum wait time for SPF calculations. Valid values range from 1 to 600000 milliseconds. <p>Note For LSA throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.</p>

Element	Description
Default Information Originate	<p>Settings used by an ASBR to generate a default external route into an OSPFv3 routing domain:</p> <ul style="list-style-type: none"> • Enable Default Information Originate – Check this box to enable generation of a default route into the OSPFv3 routing domain; the following options become available: <ul style="list-style-type: none"> • Always advertise the default route – Check this box to always advertise the default route. • Metric Value – The OSPFv3 metric used to generate the default route. Valid values range from 0 to 16777214. • Metric Type – The external link type associated with the default route advertised into the OSPFv3 routing domain. Choose 1 or 2, indicating a Type 1 or a Type 2 external route. The default value is 1. • Route Map – (Optional) Enter or Select the name of a route map object to apply. The routing process generates the default route if the route map is satisfied. <p>Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, on page 145.</p>
Passive Interface	<p>Passive routing helps control the advertisement of OSPFv3 routing information, and disables sending and receiving OSPFv3 routing updates on an interface.</p> <p>Enter or Select one or more interfaces, or interface objects, to enable passive OSPFv3 routing on those interfaces. IPv4 and IPv6 addresses are supported.</p>
Non Stop Forwarding Tab	
Note Non Stop Forwarding (NSF) is supported on ASA 9.3(1)+ only.	
Enable graceful-restart helper	<p>Enables graceful restart helper mode.</p> <p>When an ASA has NSF enabled, it is said to be NSF-capable and will operate in graceful restart mode. By default, the neighboring ASAs of the NSF-capable ASA will be NSF-aware and will operate in NSF helper mode. When the NSF-capable ASA is performing graceful restart, the helper ASAs assist in the nonstop forwarding recovery process.</p> <p>If you do not want the ASA to help the restarting neighbor with nonstop forwarding recovery, clear the Enable graceful-restart helper option.</p>
Enable Link State Advertisement	<p>Enables strict link-state advertisement (LSA) checking.</p> <p>Note When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.</p>

Element	Description
Enable graceful-restart (Use when Spanned Cluster or Failover configured)	Enables graceful restart on the ASA.
Length of graceful restart interval	(Optional) Specifies the length of the graceful restart interval, in seconds. The range is from 1 to 1800. The default is 120. Note For a restart interval below 30 seconds, graceful restart will be terminated.

Area Tab (OSPFv3)

Use the Area panel on the [Process Tab](#), on page 114 of the OSPFv3 page to configure OSPFv3 areas, ranges and virtual links. The Area panel consists of three definition tables—Area, Range, and Virtual Link:

- Refer to [Add/Edit Area Dialog Box \(OSPFv3\)](#), on page 119 for information about adding and editing Area table entries.
- Refer to [Add/Edit Range Dialog Box \(OSPFv3\)](#), on page 121 for information about adding and editing Range table entries.
- Refer to [Add/Edit Virtual Link Dialog Box \(OSPFv3\)](#), on page 122 for information about adding and editing Virtual Link table entries.

Refer to [Using Tables](#) for basic information about working with Security Manager tables.

Navigation Path

You can access the Area tab from the [Process Tab](#), on page 114 of the OSPFv3 page. For more information about the OSPFv3 page, see [Configuring OSPFv3](#), on page 112.

Related Topics

- [About OSPFv3](#), on page 112
- [OSPFv3 Interface Tab](#), on page 125

Add/Edit Area Dialog Box (OSPFv3)

Use the Add/Edit Area dialog box to define parameters for the area.

Navigation Path

You can access the Add/Edit Area dialog box from the [Area Tab \(OSPFv3\)](#), on page 119.

Related Topics

- [Configuring OSPFv3](#), on page 112
- [About OSPFv3](#), on page 112
- [Process Tab](#), on page 114

Field Reference

Table 78: Add/Edit Area Dialog Box

Element	Description
Area ID	Enter an identifier for the area as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295.
Cost	The cost of sending a packet on an interface. Valid values are 0 to 65535. Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The ASA calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.
Type	<p>Define the area type by choosing one of the following:</p> <ul style="list-style-type: none"> • Normal – Make the area a standard OSPFv3 area. This option is selected by default when you first create an area. • NSSA – Make the area a “not-so-stubby area.” NSSAs accept Type 7 LSAs. When you choose this option, the Default Information Originate options are enabled. <p>When you create a NSSA, you can prevent summary LSAs from being flooded into the area by deselecting <i>Allow sending summary LSA into this area</i> . You can also disable route redistribution by deselecting <i>Redistribute</i> , and enabling <i>Default information originate</i> .</p> <ul style="list-style-type: none"> • Stub – Make the area a stub area. Stub areas do not have any routers or areas beyond it. Stub areas prevent AS External LSAs (Type 5 LSAs) from being flooded into the stub area. When you choose this option, <i>Allow sending summary LSA into this area</i> is enabled. <p>When you create a stub area, you can prevent summary LSAs (Type 3 and 4) from being flooded into the area by deselecting <i>Allow sending summary LSA into this area</i> .</p>
<p>Default Information Originate</p> <p>These options are enabled when you choose NSSA as the area Type. The first option is enabled when you choose Stub as the area Type.</p>	
Allow sending summary LSA into this area	Select to allow flooding of summary LSAs into the area.
Redistribute (imports routes to normal and NSSA areas)	Select to allow route redistribution.

Element	Description
Default information originate	<p>Check this box to generate a Type 7 default into the NSSA. Selecting this option enables the following metric options:</p> <ul style="list-style-type: none"> • Metric – The OSPF metric value for the default route. Valid values range from 1 to 16777214. The default is 1. • Metric Type – The OSPF metric type for the default route. Choose 1 (Type 1) or 2 (Type 2). The default is 1.

Add/Edit Range Dialog Box (OSPFv3)

Use the Add/Edit Area Range Network dialog box to add a new range to the area selected in the Area table, or to change an existing entry.

Navigation Path

You can access the Add/Edit Range dialog box from the Range panel under the [Area Tab \(OSPFv3\)](#), on [page 119](#).

Related Topics

- [Configuring OSPFv3](#), on page 112
- [About OSPFv3](#), on page 112
- [Process Tab](#), on page 114

Field Reference

Table 79: Add/Edit Range Dialog Box

Element	Description
Area ID	This read-only entry is the ID of the area to which this range applies.
IPv6 Prefix/Length	<p>The IPv6 address(es) for the routes being summarized.</p> <p>Tip You can click Select to select the networks from a list of network objects.</p>
Cost	<p>The cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values are 0 to 16777215.</p> <p>Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The ASA calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.</p>
Advertise	Select this option to set the address range status to advertise. This causes Type 3 summary LSAs to be generated (this is the default). Deselect this option to suppress the Type 3 summary LSAs for the specified networks.

Add/Edit Virtual Link Dialog Box (OSPFv3)

Use the Add/Edit Virtual Link dialog box to define virtual links for the area selected in the Area table, or change the properties of existing virtual links.

Navigation Path

You can access the Add/Edit Virtual Link dialog box from the Virtual Link panel under the [Area Tab \(OSPFv3\)](#), on page 119.

Related Topics

- [Configuring OSPFv3](#), on page 112
- [About OSPFv3](#), on page 112
- [Process Tab](#), on page 114

Field Reference

Table 80: Add/Edit Virtual Link Dialog Box

Element	Description
Area ID	This read-only entry is the ID of the area to which this virtual link applies.
Peer Router ID	Enter the IP address of the virtual link neighbor. Tip You can click Select to select from a list of network objects.
TTL Security	The time-to-live (TTL) security hop count on a virtual link. The hop count value can range from 1 to 254.
Dead Interval	The interval, in seconds, if no hello packets are received, neighbors declare the device down. Valid values range from 1 to 8192. The default value of this field is four times the Hello Interval.
Hello Interval	The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 8192 seconds. The default value is 10 seconds.
Transmit Interval	The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a device sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the device does not receive an acknowledgment, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 8192 seconds. The default value is 5 seconds.

Element	Description
Transmit Delay	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 8192 seconds. The default value is 1 second.

Add/Edit Redistribution Dialog Box (OSPFv3)

Use the Add/Edit Redistribution dialog box to add a redistribution rule to this process, or to edit an existing redistribution rule.

Navigation Path

You can access the Redistribution dialog box from the Redistribution panel under the [Process Tab](#), on page 114.

Related Topics

- [Configuring OSPFv3](#), on page 112
- [About OSPFv3](#), on page 112

Field Reference

Table 81: Add/Edit Redistribution Dialog Box

Element	Description
Source Protocol	Choose the source protocol for route redistribution: <ul style="list-style-type: none"> • Connected – Redistributes connected routes (routes established automatically by virtue of having an IP address enabled on the interface) to the OSPFv3 routing process. Connected routes are redistributed as external to the autonomous system. • OSPF – Redistributes routes from another OSPF routing process. The Routing PID and the Match options are enabled when you choose this option. • Static – Redistributes static routes to the OSPFv3 routing process.
Metric	The metric value for the routes being redistributed. Valid values range from 1 to 16777214; the default is 20. When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified.

Element	Description
Metric Type	The metric type is the external link type associated with the default route that is advertised into the OSPFv3 routing domain. Choose None, 1, or 2, where None indicates there is no default route, 1 indicates the metric is a Type 1 external route, and 2 is a Type 2 external route.
Tag (optional)	The tag is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between other border devices. Valid values range from 0 to 4294967295.
Route Map	Enter or Select the name of the route map object to apply to the redistribution entry. Tip Click Select to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects , on page 145.
Routing PID	The ID of the process to which redistribution is directed. (The Process ID is defined on the Process Tab , on page 114.) This option is enabled only when OSPF is chosen as the Source Protocol.
Include Connected	Check this box to include connected routes in the redistribution.
Match	
The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions. These options are enabled only when OSPF is chosen as the Source Protocol.	
Internal	The route is internal to a specific autonomous system.
External 1	Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
External 2	Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
NSSA External 1	Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.
NSSA External 2	Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.

Add/Edit Summary Prefix Dialog Box (OSPFv3)

Use the Add/Edit Summary Prefix dialog box to add new route-summarization entries to the selected process, or to modify existing entries.

Navigation Path

You can access the Add/Edit Summary Prefix dialog box from the Summary Prefix panel under the [Process Tab](#) , on page 114.

Related Topics

- [Configuring OSPFv3](#) , on page 112
- [About OSPFv3](#) , on page 112

Field Reference

Table 82: Add/Edit Summary Prefix Dialog Box

Element	Description
Process ID	This read-only value identifies the process to which this rule applies.
IPv6 Prefix/Length	Enter an IPv6 prefix/length for external route summarization. Tip You can click Select to select from a list of network objects.
Advertise	When selected, summary routes that match the specified prefix and mask pair are advertised. When deselected, routes that match the specified prefix and mask pair are suppressed. By default, this check box is selected.
Tag (optional)	The tag is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between border devices. Valid values range from 0 to 4294967295. This field is enabled when you check Advertise.

OSPFv3 Interface Tab

Use the Interface panel to configure interface-and neighbor-specific OSPFv3 routing properties. The Interface panel consists of two definition tables, Interface and Neighbor:

- Refer to [Add/Edit Interface Dialog Box \(OSPFv3\)](#) , on page 126 for information about adding and editing Interface table entries.
- Refer to [Add/Edit Neighbor Dialog Box \(OSPFv3\)](#) , on page 129 for information about adding and editing Neighbor table entries.

Refer to [Using Tables](#) for basic information about working with Security Manager tables.

Navigation Path

Click the Interface tab on the OSPFv3 page to display this panel. For more information about the OSPFv3 page, see [Configuring OSPFv3](#) , on page 112.

Related Topics

- [About OSPFv3](#) , on page 112

- [Process Tab](#) , on page 114

Add/Edit Interface Dialog Box (OSPFv3)

Use the Add/Edit Interface dialog box to define OSPFv3 routing properties for an individual interface, or to change an existing entry.

Navigation Path

You can access the Add/Edit Interface dialog box from the Interface panel under the [OSPFv3 Interface Tab](#) , on page 125.

Related Topics

- [Configuring OSPFv3](#) , on page 112
- [About OSPFv3](#) , on page 112
- [Process Tab](#) , on page 114

Field Reference

Table 83: Add/Edit Interface Dialog Box

Element	Description
Interface	The name of the interface to which this routing configuration applies. Tip You can click Select to select from a list of interface objects.
Enable OSPFv3 on this interface	Check this box to enable OSPFv3 on the specified interface, and activate the following fields: <ul style="list-style-type: none"> • Process ID – Choose the process to apply to this interface; defined on the OSPFv3 Process Tab , on page 114. • Area ID – Identify the area to be assigned; areas are also defined on the OSPFv3 Process Tab , on page 114. • Instance ID – (Optional) Specify an ID for this process instance. Valid values for this setting range from 0 to 255. <p>This feature lets you have multiple OSPFv3 processes on a single link. Received packets with other instance IDs are then ignored by this process.</p>
Properties	
Filter outgoing link-state advertisements	Check this box to filter outgoing LSAs. All outgoing LSAs are flooded to the interface by default.
Disable MTU mismatch detection	Check this box to disable the OSPFv3 MTU mismatch detection when database description (DBD) packets are received.
Flood Reduction	Check this box to suppress unnecessary flooding of LSAs in stable topologies.

Element	Description
Point-to-point Network	<p>Check this box to define this as a link to a point-to-point network; that is, a network between two routing devices. All neighbors on a point-to-point network establish adjacency and there is no designated router.</p> <p>This option is unavailable when the Broadcast option is selected.</p>
Broadcast	<p>Check this box to define this as a link to a network with multiple routing devices. Such networks establish a designated router (DR), as well as a backup designated router (BDR), that controls LSA flooding on the network.</p> <p>This option is unavailable when the Point-to-point Network option is selected.</p>
Cost	<p>The cost of sending a packet through the interface. Link cost is an arbitrary number used in shortest path first calculations. If you do not assign a value, the configured reference bandwidth divided by the interface port speed is used. (The default reference bandwidth is 40 Gb/sec.)</p>
Priority	<p>Assign an OSPFv3 priority to this interface. Valid values for this setting range from 0 to 255. Entering 0 for this setting makes the device ineligible to become the designated router or backup designated router. This setting does not apply to interfaces that are configured as point-to-point, non-broadcast interfaces.</p> <p>When two routing devices connect to a network, both attempt to become the designated router. The device with the higher priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router.</p>
Dead Interval	<p>If no hello packets are received from a neighbor within this interval, that device is designated as inactive. Valid values range from 1 to 65535. The default value for this setting is four times the hello interval.</p>
Poll Interval	<p>If a neighboring device is inactive, it may be necessary to continue sending hello packets to that neighbor. The hello packets are sent at this reduced interval, which should be larger than the hello interval.</p>
Retransmit Interval	<p>The time, in seconds, between LSA retransmissions for adjacent neighbors. When a router sends an LSA to a neighbor, it keeps the LSA until it receives an acknowledgment. If an acknowledgment is not received within this interval, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds.</p>
Transmit Delay	<p>The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds.</p>
Authentication	

Element	Description
Type	<p>The type of authentication enabled on this interface. Choose one of the following:</p> <ul style="list-style-type: none"> • Area – OSPFv3 does not provide “built-in” authentication, instead relying on IPv6/IPSec protocols. Choose this option to use those protocols to authenticate OSPFv3 traffic on all interfaces in the area; this means all routing devices in the area must use this option. This is the default. • Interface – Choose this option to secure this interface and protect OSPFv3 virtual links. The additional parameters in this section are enabled when you choose this option. • None – OSPFv3 authentication is disabled.
Security Parameter Index	<p>Enter an IPSec identification tag used to distinguish this particular OSPFv3 interface; used in conjunction with the specified authentication and encryption rules. Valid values range from 256 to 4294967295.</p>
Authentication Algorithm	<p>Choose the type of authentication algorithm to use:</p> <ul style="list-style-type: none"> • md5 – Message Digest 5; produces a 128-bit hash value. • sha1 – Secure Hash Algorithm version 1; produces a 160-bit hash value.
Authentication Key	<p>Enter an authentication key. The length of the key entered depends on the type of authentication chosen as the Authentication Algorithm, and whether the key is to be encrypted (when you check the Encrypt Authentication Key box):</p> <ul style="list-style-type: none"> • md5 – 32 characters. • md5 (encrypted) – 66 characters. • sha1 – 40 characters. • sha1 (encrypted) – 82 characters.
Encrypt Authentication Key	<p>Check this box to require encryption of the specified Authentication Key for transmission.</p>
Include Encryption	<p>Check this box to require encryption of OSPFv3 packets. The following options are enabled.</p>

Element	Description
Encryption Algorithm	<p>Choose the type of encryption to use:</p> <ul style="list-style-type: none"> • 3des – Triple DES; the Data Encryption Standard cipher algorithm is applied three times to each packet. • aes-cbc – Encryption is based on the Advanced Encryption Standard with Cipher Block Chaining, to produce a key of the size chosen with the Key Type parameter. <p>The Key Type list is enabled only when you choose this encryption option. Choose one of these options:</p> <ul style="list-style-type: none"> • 128 – For 128-bit keys. • 192 – For 192-bit keys. • 256 – For 256-bit keys. • des – Encryption is based on the Data Encryption Standard, using 56-bit keys.
Encryption Key	<p>Enter an encryption key. The length of the key entered depends on the type of encryption chosen as the Encryption Algorithm, and whether the key is to be encrypted (when you check the Encrypt Key box):</p> <ul style="list-style-type: none"> • 3des – 48 characters (192 bits). • 3des (encrypted) – 98 characters (192 bits). • aes-cbc/128 – 32 characters (128 bits). • aes-cbc/128 (encrypted) – 66 characters (128 bits). • aes-cbc/192 – 48 characters (192 bits). • aes-cbc/192 (encrypted) – 98 characters (192 bits). • aes-cbc/256 – 64 characters (256 bits). • aes-cbc/256 (encrypted) – 130 characters (256 bits). • des – 16 characters (64 bits). • des (encrypted) – 34 characters (64 bits).
Encrypt Key	Check this box to require encryption of the specified Encryption Key for transmission.

Add/Edit Neighbor Dialog Box (OSPFv3)

You must define a static neighbor for each point-to-point, non-broadcast interface. This feature lets you broadcast OSPFv3 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel. Note the following restrictions:

- You cannot define the same static neighbor for two different OSPFv3 processes.
- You must define a static route for each static neighbor.

Use the Add/Edit Neighbor dialog box to define a static neighbor for the interface selected in the Interface table, or to change information for an existing static neighbor.

Navigation Path

You can access the Add/Edit Neighbor dialog box from the Neighbor panel under the [OSPFv3 Interface Tab](#), on page 125.

Related Topics

- [Configuring OSPFv3](#), on page 112
- [About OSPFv3](#), on page 112
- [Process Tab](#), on page 114

Field Reference

Table 84: Add/Edit Neighbor Dialog Box

Element	Description
Interface	The interface associated with this neighbor definition (read-only).
Link-local Address	Enter the IPv6 address of the static neighbor.
Cost and Database Filter	<p>Check this box to enable filtering of the outgoing LSAs on the interface during synchronization and flooding. The following options are enabled:</p> <ul style="list-style-type: none"> • Cost – Use this field to assign an arbitrary cost to this neighbor. If a value is not assigned, the cost of the interface is used (this value is based on the port speed of the interface, and is calculated as reference bandwidth divided by interface speed). Valid values range from 1 to 65535. • Filter outgoing link-state advertisements – Check this box to disable forwarding of outgoing LSAs to this neighbor. <p>Note The Cost and Database Filter options and the Poll-Interval options are mutually exclusive.</p>

Element	Description
Poll-Interval	<p>Check this box to enable the following options:</p> <ul style="list-style-type: none"> • Poll Interval – Time interval in seconds between transmission of hello packets to a “dead” neighbor. The default is 120. <p>If a neighboring device becomes inactive (hello packets have not been received for the dead interval period), it may be necessary to continue sending hello packets to the dead neighbor at a reduced rate. Thus this value should be larger than the interface hello interval.</p> <ul style="list-style-type: none"> • Priority – The router priority value of this neighbor. The default is 0; valid values range from 1 to 255. <p>The priority value helps determine the designated router for an OSPFv3 link. A value of zero means the device is ineligible to become the designated router, or backup designated router.</p> <p>Note The Poll-Interval options and the Cost and Database Filter options are mutually exclusive. Also, these values do not apply to point-to-multipoint interfaces.</p>

Configuring RIP

Routing Information Protocol (RIP) is a dynamic routing protocol, or more precisely, an interior gateway protocol that is based on distance vectors. RIP uses hop count as the metric for path selection. When RIP is enabled on an interface, the interface exchanges RIP broadcast packets with neighboring devices to dynamically learn about and advertise routes. These RIP packets contain information about the destination networks that the gateways can reach, and the number of gateways that a packet must travel through to reach those destinations.

Cisco Security Manager supports both RIP version 1 and RIP version 2. Version 1 does not send the subnet mask with the routing update; RIP version 2 sends the subnet mask with the routing update, and supports variable-length subnet masks. Additionally, RIP version 2 supports neighbor authentication when routing updates are exchanged. This authentication ensures that the security appliance receives reliable routing information from a trusted source.



Note You cannot enable RIP if you have OSPF processes running.

Limitations

RIP has the following limitations:

- Cisco Security Manager cannot pass RIP updates between interfaces.
- RIP Version 1 does not support variable-length subnet masks.
- RIP has a maximum hop count of 15. A route with a hop count greater than 15 is considered unreachable.
- RIP convergence is relatively slow compared to other routing protocols.

RIP Version 2 Notes

The following information applies to RIP Version 2 only:

- If using neighbor authentication, the authentication key and key ID must be the same on all neighbor devices that provide RIP version 2 updates to the interface.
- With RIP version 2, the security appliance transmits and receives default route updates using the multicast address 224.0.0.9. In passive mode, it receives route updates at that address.
- When RIP version 2 is configured on an interface, the multicast address 224.0.0.9 is registered on that interface. When a RIP version 2 configuration is removed from an interface, that multicast address is unregistered.

Using Security Manager to Configure RIP on Security Appliances

Use the RIP page to enable the Routing Information Protocol on an interface. The settings and features available when configuring RIP depend on the type of device and OS version that you are configuring:

- To configure RIP on a PIX Firewall or ASA running an OS version earlier than 7.2, or on any FWSM, see [RIP Page for PIX/ASA 6.3–7.1 and FWSM](#), on page 132.
- To configure RIP on a PIX Firewall or ASA running OS version 7.2 or later, see [RIP Page for PIX/ASA 7.2 and Later](#), on page 134.

Related Topics

- [Configuring Static Routes](#), on page 141
- [Configuring OSPF](#), on page 80
- [Configuring No Proxy ARP](#), on page 1
- [Configuring Routing Information Protocol](#) – a chapter from the “Cisco IOS IP Configuration Guide, Release 12.2,” providing additional detailed information about RIP

RIP Page for PIX/ASA 6.3–7.1 and FWSM



Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any bug fixes or enhancements.

Use this RIP page to enable the Routing Information Protocol (RIP) on an interface in any FWSM, or in a PIX/ASA running a pre-7.2 version operating system.

The RIP table on this page lists all interfaces on which RIP is currently defined. Use the Add RIP Configuration and Edit RIP Configuration dialog boxes to create and maintain these entries. See [RIP Page for PIX/ASA 6.3–7.1 and FWSM](#), on page 132 for more information.

Navigation Path

- (Device view) Select **Platform > Routing > RIP** from the Device Policy selector.

- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > RIP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

When creating a shared RIP policy, you must choose a Version in the Create a Policy dialog box, as follows:

- **PIX/ASA 6.3-7.1 and FWSM**
- **PIX/ASA 7.2 and Later**

When assigning a shared RIP policy, be sure to assign the appropriate RIP policy for the device. For example, you cannot assign a PIX/ASA 7.2+ RIP policy to an FWSM.

Related Topics

- [Configuring Static Routes](#) , on page 141
- [Configuring OSPF](#) , on page 80
- [Configuring No Proxy ARP](#) , on page 1
- [RIP Page for PIX/ASA 7.2 and Later](#) , on page 134
- Standard rules table topics:
 - [Using Rules Tables](#)
 - [Table Columns and Column Heading Features](#)

Add/Edit RIP Configuration (PIX/ASA 6.3–7.1 and FWSM) Dialog Boxes



Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any bug fixes or enhancements.

Use the Add RIP Configuration and Edit RIP Configuration dialog boxes to add a RIP configuration to the security appliance, or to make changes to an existing RIP configuration. By adding a RIP configuration, you enable RIP on the specified interface. Except for their titles, the two dialog boxes are identical.

Navigation Path

You can access the Add and Edit RIP Configuration dialog boxes from the [RIP Page for PIX/ASA 6.3–7.1 and FWSM](#) , on page 132.

Field Reference

Table 85: Add/Edit RIP Configuration (PIX/ASA 6.3-7.1 and FWSM) Dialog Boxes

Element	Description
Interface	Enter or Select the interface for the RIP configuration. You cannot configure two different RIP configurations on the same interface.

Element	Description
Mode	<p>Select the interface behavior regarding RIP updates:</p> <ul style="list-style-type: none"> • Send default routes – The interface will transmit RIP routing updates only. • Receive routes – The interface will listen for RIP routing broadcasts and use that information to populate its routing table, but it will not send RIP routing updates. • Send default routes and receive routes – The interface will send and receive RIP routing updates.
Version	<p>Select the RIP version to enable on the interface:</p> <ul style="list-style-type: none"> • RIP Version 1 – Enables RIP Version 1 on the interface. • RIP Version 2 – Enables RIP Version 2 on the interface. Configuring RIP Version 2 registers the multicast address 224.0.0.9 on the interface.
Version 2 Authentication	<p>These options let you enable and select the type of authentication used with RIP Version 2.</p> <ul style="list-style-type: none"> • Enable Authentication – This option is available when you select RIP Version 2 above. When this box is checked, RIP neighbor authentication is enabled and the following options become available: <ul style="list-style-type: none"> • Type – Select MD5 to use the MD5 hash algorithm for authentication (recommended), or select Clear text to use clear text for authentication. • Key ID – The identification number of the authentication key. This number must be shared with all other devices sending updates to and receiving updates from the security appliance. Valid values range from 1 to 255. • Key – The shared key used for authentication. This key must be shared with all other devices sending updates to and receiving updates from the security appliance. The key can be up to 16 characters.

RIP Page for PIX/ASA 7.2 and Later



Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any bug fixes or enhancements.

Use this RIP page to enable and configure the Routing Information Protocol (RIP) on PIX and ASA devices running operating system 7.2 or later. The RIP page consists of these tabbed panels:

- [RIP - Setup Tab , on page 135](#)
- [RIP - Redistribution Tab , on page 137](#)
- [RIP - Filtering Tab , on page 138](#)
- [RIP - Interface Tab , on page 139](#)

Navigation Path

- (Device view) Select **Platform > Routing > RIP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > RIP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

When creating a shared RIP policy, you must choose a Version in the Create a Policy dialog box, as follows:

- **PIX/ASA 6.3-7.1 and FWSM**
- **PIX/ASA 7.2 and Later**

When assigning a shared RIP policy, be sure to assign the appropriate RIP policy for the device. For example, you cannot assign a PIX/ASA 7.2+ RIP policy to an FWSM.

Related Topics

- [Configuring Static Routes](#) , on page 141
- [Configuring OSPF](#) , on page 80
- [Configuring No Proxy ARP](#) , on page 1
- [RIP Page for PIX/ASA 6.3–7.1 and FWSM](#) , on page 132

RIP - Setup Tab

Use the Setup panel to define RIP on the security appliance, and to configure global RIP protocol parameters. You can only enable a single RIP process on the security appliance.

Navigation Path

You can access the Setup tab from the [RIP Page for PIX/ASA 7.2 and Later](#) , on page 134.

Related Topics

- [RIP - Redistribution Tab](#) , on page 137
- [RIP - Filtering Tab](#) , on page 138
- [RIP - Interface Tab](#) , on page 139

Field Reference

Table 86: Setup Tab

Element	Description
Networks	<p>Define one or more networks for RIP routing. Enter IP address(es), or enter or Select the desired Network/Hosts objects (see Understanding Networks/Hosts Objects); IP addresses must not contain any subnet information. There is no limit to the number of networks you can add to the security appliance configuration.</p> <p>The RIP routing updates will be sent and received only through interfaces on the specified networks. Also, if the network of an interface is not specified, the interface will not be advertised in any RIP updates.</p>
Passive Interface	<p>Use this option to specify passive interfaces on the security appliance, and by extension the active interfaces. The device listens for RIP routing broadcasts on passive interfaces, using that information to populate its routing tables, but does not broadcast routing updates on passive interfaces. Interfaces that are not designated as passive, receive and send updates. Choose one of these options:</p> <ol style="list-style-type: none"> None – No interfaces are designated as passive. All Interfaces – All interfaces on the device are designated as passive, except those entered the Excluded Interfaces field below. Specified Interfaces – Only those interfaces explicitly specified in the Interfaces field below are designated as passive.
Interfaces/Excluded Interfaces	<p>Use this field to specify the interfaces excluded from the passive list, or those explicitly designated as passive, depending on your choice from the Passive Interface list above:</p> <ul style="list-style-type: none"> If you chose All Interfaces, this field is labeled Excluded Interfaces: enter or Select only those interfaces to be excluded (that is, those that are to be active not passive). If you chose Specified Interfaces in the Passive Interface list, enter or Select those interfaces that are to be designated as passive. <p>Note You cannot specify two different RIP configurations for the same interface.</p>
RIP Version	<p>Choose the RIP versions for sending and receiving RIP updates:</p> <ul style="list-style-type: none"> Receive Version 1 and 2, Send Version 1 Send and Receive Version 1 Send and Receive Version 2
Generate Default Route	<p>When selected, a default route is generated for distribution, based on the Route Map you specify.</p>

Element	Description
Route Map	Specify the route map to use for generating default routes. Note This field contains only the Route Map name. The Route Map is created and contained within a FlexConfig; see Understanding FlexConfig Policies and Policy Objects for more information.
Enable Auto-Summary	When Send and Receive Version 2 is the chosen RIP Version, this option is available. When checked, automatic route summarization is enabled. Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is disabled, subnets are advertised. Note RIP Version 1 always uses automatic summarization—you cannot disable it.

RIP - Redistribution Tab

Use the Redistribution panel to manage redistribution routes. These are the routes that are being redistributed from other routing processes into the RIP routing process. See [Add/Edit Redistribution Dialog Box](#), on page 32 for more information.

Navigation Path

You can access the Redistribution tab from the [RIP Page for PIX/ASA 7.2 and Later](#), on page 134.

Related Topics

- [RIP - Setup Tab](#), on page 135
- [RIP - Filtering Tab](#), on page 138
- [RIP - Interface Tab](#), on page 139

Add/Edit Redistribution Dialog Box

Use the Add Redistribution and Edit Redistribution dialog boxes to add and edit redistribution routes on the [RIP - Redistribution Tab](#), on page 137. These are the routes that are being redistributed from other routing processes into the RIP routing process. Except for their titles, these two dialog boxes are identical.

Navigation Path

You can access the Add and Edit Redistribution dialog boxes from the Redistribution tab on the [RIP Page for PIX/ASA 7.2 and Later](#), on page 134.

Field Reference

Table 87: Add/Edit Redistribution Dialog Box

Element	Description
Protocol to Redistribute	<p>Choose the routing protocol to redistribute into the RIP routing process:</p> <ul style="list-style-type: none"> • Static - Static routes • Connected – Directly connected networks. • OSPF – Routes discovered by the OSPF routing process. <p>If you choose OSPF, you must also enter the OSPF Process ID and, optionally, Match criteria.</p>
Process ID	Enter the process ID when the OSPF protocol is chosen.
Match	<p>If you are redistributing OSPF routes into the RIP routing process, you can select specific types of OSPF routes to redistribute. Ctrl-click to select multiple types:</p> <ul style="list-style-type: none"> • Internal – Routes internal to the autonomous system (AS) are redistributed. • External 1 – Type 1 routes external to the AS are redistributed. • External 2 – Type 2 routes external to the AS are redistributed. • NSSA External 1 – Type 1 routes external to a not-so-stubby area (NSSA) are redistributed. • NSSA External 2 – Type 2 routes external to an NSSA are redistributed. <p>Match criteria are optional. The default is match Internal, External 1, and External 2.</p>
Metric	<p>The RIP metric type to apply to the redistributed routes. The two choices are:</p> <ul style="list-style-type: none"> • Transparent – Use the current route metric. • Specified Value – Assign a specific metric value.
Route Map	<p>The name of a route map that must be satisfied before the route can be redistributed into the RIP routing process.</p> <p>Note This field contains only the route Map name. The contents of the route map are created and contained within a FlexConfig. See Understanding FlexConfig Policies and Policy Objects for more information.</p>

RIP - Filtering Tab

Use the Filtering panel to manage filters for the RIP policy. Filters are used to limit network information in incoming and outgoing RIP advertisements. See [Add/Edit Filter Dialog Box](#), on page 139 for more information.

Navigation Path

You can access the Filtering tab from the [RIP Page for PIX/ASA 7.2 and Later](#), on page 134.

Related Topics

- [RIP - Setup Tab , on page 135](#)
- [RIP - Redistribution Tab , on page 137](#)
- [RIP - Interface Tab , on page 139](#)

Add/Edit Filter Dialog Box

Use the Add Filter and Edit Filter dialog boxes to add and edit RIP filters on the [RIP - Filtering Tab , on page 138](#). Filters are used to limit network information in incoming and outgoing RIP advertisements. Except for their titles, these two dialog boxes are identical.

Navigation Path

You can access the Add and Edit Filter dialog boxes from the Filtering tab on the [RIP Page for PIX/ASA 7.2 and Later , on page 134](#).

Field Reference

Table 88: Add/Edit Filter Dialog Box

Element	Description
Traffic Direction	Choose the type of traffic to be filtered: Inbound or Outbound . Note If Traffic Direction is Inbound, you can define an Interface filter only.
Filter On	Specify whether the filter is based on an Interface or a Route . If you select Interface, enter or Select the name of the interface on which routing updates are to be filtered. If you select Route, choose the route type: <ul style="list-style-type: none"> • Static – Only static routes are filtered. • Connected – Only connected routes are filtered. • OSPF – Only OSPF routes discovered by the specified OSPF process are filtered. Enter the Process ID of the OSPF process to be filtered.
Filter ACLs	Enter or Select the name of one or more access control lists (ACLs) that define the networks to be allowed or removed from RIP route advertisements.

RIP - Interface Tab

Use the Interface panel to manage the interfaces configured to send and receive RIP broadcasts. See [Add/Edit Interface Dialog Box , on page 140](#) for more information.

Navigation Path

You can access the Interface tab from the [RIP Page for PIX/ASA 7.2 and Later , on page 134](#).

Related Topics

- [RIP - Setup Tab](#) , on page 135
- [RIP - Redistribution Tab](#) , on page 137
- [RIP - Filtering Tab](#) , on page 138

Add/Edit Interface Dialog Box

Use the Add Interface and Edit Interface dialog boxes to add and edit RIP interface configurations on the [RIP - Interface Tab](#) , on page 139. Except for their titles, these two dialog boxes are identical.

Navigation Path

You can access the Add and Edit Interface dialog boxes from the Interface tab on the [RIP Page for PIX/ASA 7.2 and Later](#) , on page 134.

Field Reference

Table 89: Add/Edit Interface Dialog Box

Element	Description
Interface	Enter or Select an interface defined on this appliance.
Send (Version)	These options let you override, for this interface, the global Send versions specified on the RIP - Setup Tab , on page 135. Select the appropriate boxes to specify sending updates using RIP Version 1, Version 2, or both.
Receive (Version)	These options let you override the global Receive versions. Select the appropriate boxes to specify accepting updates using RIP Version 1 only, Version 2 only, or both.
Authentication Type	<p>Choose the authentication used on this interface for RIP broadcasts:</p> <ul style="list-style-type: none"> • None – No authentication. • MD5 – Employ MD5. • Clear Text – Employ clear-text authentication. <p>If you choose MD5 or Clear Text, you must also provide the following authentication parameters:</p> <ul style="list-style-type: none"> • Key ID – The ID of the authentication key. Valid values are from 0 to 255. • Key – The key used by the chosen authentication method. Can contain up to 16 characters. • Confirm – Enter the authentication key again, to confirm.

Configuring Static Routes

A static route is a specific path to a particular destination network that is manually defined on the current device. Static routes are used in a variety of situations, and can be a quick and effective way to route data from one network to another when there is no dynamic route to the destination, or when use of a dynamic routing protocol is not feasible.

All routes have a value or “metric” that represents its priority of use. (This metric is also referred to as “administrative distance.”) When two or more routes to the same destination are available, devices use administrative distance to decide which route to use.

For static routes, the default metric value is one, which gives them precedence over routes from dynamic routing protocols. If you increase the metric to a value greater than that of a dynamic route, the static route operates as a back-up in the event that dynamic routing fails. For example, Open Shortest Path First (OSPF)-derived routes have a default administrative distance of 100. To configure a back-up static route that is overridden by an OSPF route, specify a metric value for the static route that is greater than 100. This is referred to as a “floating” static route.

There is a special kind of static route known as a default route, or a “zero-zero” route because all zeroes are used for both the destination address and subnet mask. The default static route serves as a catch-all gateway: if there are no matches for a particular destination in the device’s routing table, the default route is used. The default route generally includes a next-hop IP address or local exit interface.

Use the Static Route page to maintain manually defined static routes. The Static Route table on this page lists all currently defined static routes, showing for each, the name of the interface or interface role for which the route is defined, the destination network(s), the next hop gateway, the route metric, whether the route is tunneled, and whether there is service-level agreement tracking for the route. For a detailed explanation of these fields, see [Add/Edit Static Route Dialog Box](#), on page 142 or [Add/Edit IPv6 Static Route Dialog Box](#), on page 143.

Static null0 Route Configuration

Typically ACLs are used for traffic filtering and they enable you to filter packets based on the information contained in their headers. In packet filtering, the ASA firewall examines packet headers to make a filtering decision, thus adding some overhead to the processing of the packets and affecting performance.

Static null 0 routing is a complementary solution to filtering. A static null0 route is used to forward unwanted or undesirable traffic into a black hole. The null interface null0, is used to create the black hole. Static routes are created for destinations that are not desirable, and the static route configuration points to the null interface. Any traffic that has a destination address that has a best match of the black hole static route is automatically dropped. Unlike with ACLs, static null0 routes do not cause any performance degradation.

The static null0 route configuration is used to prevent routing loops. BGP leverages the static null0 configuration for Remotely Triggered Black Hole routing.

Navigation Path

- (Device view) Select **Platform > Routing > Static Route** or **Platform > Routing > IPv6 Static Route** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > Static Route** or **PIX/ASA/FWSM Platform > Routing > IPv6 Static Route** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- [Add/Edit Static Route Dialog Box](#) , on page 142
- [Add/Edit IPv6 Static Route Dialog Box](#) , on page 143
- [Monitoring Service Level Agreements \(SLAs\) To Maintain Connectivity](#)
- Standard rules table topics:
 - [Using Rules Tables](#)
 - [Table Columns and Column Heading Features](#)

Add/Edit Static Route Dialog Box

The Add/Edit Static Route dialog box lets you add or edit a static route.

Navigation Path

You can access the Add/Edit Static Route dialog box from the Static Routes page. Click the Add Row button to add a new static route; select an existing static route and click the Edit Row button to edit that route.

Related Topics

- [Configuring Static Routes](#) , on page 141

Field Reference

Table 90: Add/Edit Static Route Dialog Box

Element	Description
Interface	<p>Enter or Select the interface to which this static route applies.</p> <p>Sending traffic to a Null0 interface results in dropping the packets destined to the specified network. This feature is useful in configuring Remotely Triggered Black Hole (RTBH) for BGP. For more information, see Configuring Static Routes , on page 141.</p> <p>Note If Null0 is selected as the interface, the Gateway and Tunneled options are disabled.</p>
Network	<p>Enter or Select the destination network(s). You can provide one or more IP address/netmask entries, one or more Networks/Hosts objects, or a combination of both; separate the entries with commas.</p> <p>Enter “0.0.0.0/0” or “any” to specify a default route.</p>
Gateway	<p>Enter or Select the gateway router which is the next hop for this route. You can provide an IP address, or a Networks/Hosts object.</p> <p>Note If an IP address from one of the security appliance’s interfaces is used as the Gateway IP address, the security appliance will resolve the designated IP address in the packet instead of resolving the Gateway IP address.</p>

Element	Description
Metric	<p>The Metric is a measurement of the “expense” of a route, based on the number of hops (hop count) to the network on which a specific host resides. Hop count is the number of networks that a network packet must traverse, including the destination network, before it reaches its final destination. Because the hop count includes the destination network, all directly connected networks have a metric of 1.</p> <p>Enter the number of hops to the destination network. Valid values range from 1 to 255; the default value is 1.</p> <p>The maximum number of equal-cost (equal-metric) routes that can be defined per interface is three. You cannot add a route with the same metric on different interfaces that are on the same network.</p>
Tunneled	Select this option to make this a tunnel route; can be used only for a default route. You can configure only one default tunneled gateway per device. The Tunneled option is not supported in transparent mode. Available only on PIX/ASA 7.0+ devices.
Route Tracking	<p>To monitor route availability, enter or Select name of an SLA (service level agreement) object that defines the monitoring policy. Available only on PIX/ASA 7.2+ devices.</p> <p>For more information on route tracking, see Monitoring Service Level Agreements (SLAs) To Maintain Connectivity.</p>

Add/Edit IPv6 Static Route Dialog Box

The Add/Edit IPv6 Static Route dialog box lets you add or edit an IPv6 static route. IPv6 static routes are only supported on the following devices:

- ASA 7.0 and later (Routed mode)
- ASA 8.2 and later (Transparent mode)
- FWSM 3.1 and later (Routed mode)

Navigation Path

You can access the Add/Edit IPv6 Static Route dialog box from the IPv6 Static Route page. Click the **Add Row** button to add a new static route; select an existing static route and click the **Edit Row** button to edit that route.

Related Topics

- [Configuring Static Routes](#) , on page 141

Field Reference

Table 91: Add/Edit IPv6 Static Route Dialog Box

Element	Description
Interface	Enter or Select the interface to which this static route applies.

Element	Description
IPv6 Network	<p>Enter or Select the destination network(s). You can provide one or more IP address entries, one or more Networks/Hosts objects, or a combination of both; separate the entries with commas.</p> <p>Enter two colons (::) to specify a default route.</p>
IPv6 Gateway	<p>Enter or Select the gateway router which is the next hop for this route. You can provide an IP address, or a Networks/Hosts object.</p> <p>Note If an IP address from one of the security appliance's interfaces is used as the Gateway IP address, the security appliance will resolve the designated IP address in the packet instead of resolving the Gateway IP address.</p>
Metric	<p>The Metric is a measurement of the "expense" of a route, based on the number of hops (hop count) to the network on which a specific host resides. Hop count is the number of networks that a network packet must traverse, including the destination network, before it reaches its final destination. Because the hop count includes the destination network, all directly connected networks have a metric of 1.</p> <p>Enter the number of hops to the destination network. Valid values range from 1 to 255; the default value is 1.</p> <p>The maximum number of equal-cost (equal-metric) routes that can be defined per interface is three. You cannot add a route with the same metric on different interfaces that are on the same network.</p>
Tunneled	<p>Select this option to specify the route as the default tunnel gateway for VPN traffic. You can configure only one default tunneled gateway per device. Available only on ASA 7.0+ devices in routed mode.</p>

Configuring Policy Objects for ASA Routing Policies

There are several policy objects that you use with ASA routing policies. This reference explains the configuration of these policy objects.

This section contains the following topics:

- [Understanding Route Map Objects](#) , on page 145
- [Add or Edit Policy List Object Dialog Box](#) , on page 156
- [Add or Edit Prefix List Object Dialog Box](#) , on page 159
- [Add or Edit Prefix List IPv6 Object Dialog Box](#) , on page 161
- [Add or Edit As Path Object Dialog Boxes](#) , on page 163
- [Add or Edit Community List Object Dialog Box](#) , on page 165
- [Create BFD Template](#), on page 76

Understanding Route Map Objects

You can use route maps to define the conditions for redistributing routes from one routing protocol into another, or to enable policy routing.

Route maps have many features in common with widely known ACLs. These are some of the traits common to both:

- They are an ordered sequence of individual statements, each has a permit or deny result. Evaluation of ACL or route maps consists of a list scan, in a predetermined order, and an evaluation of the criteria of each statement that matches. A list scan is aborted once the first statement match is found and an action associated with the statement match is performed.
- They are generic mechanisms—Criteria matches and match interpretation are dictated by the way that they are applied. The same route map applied to different tasks might be interpreted differently.

These are some of the differences between route maps and ACLs:

- Route maps frequently use ACLs as matching criteria.



Note Route maps do not support ACLs that include a user, user group, security group tag, or fully qualified domain name objects.

- The main result from the evaluation of an ACL is a yes or no answer—An ACL either permits or denies input data. Applied to redistribution, an ACL determines if a particular route can (route matches ACLs permit statement) or can not (matches deny statement) be redistributed. Typical route maps not only permit (some) redistributed routes but also modify information associated with the route, when it is redistributed into another protocol.
- Route maps are more flexible than ACLs and can verify routes based on criteria which ACLs can not verify. For example, a route map can verify if the type of route is internal.
- Each ACL ends with an implicit deny statement, by design convention; there is no similar convention for route maps. If the end of a route map is reached during matching attempts, the result depends on the specific application of the route map. Fortunately, route maps that are applied to redistribution behave the same way as ACLs: if the route does not match any clause in a route map then the route redistribution is denied, as if the route map contained deny statement at the end.

Route maps are preferred if you intend to either modify route information during redistribution or if you need more powerful matching capability than an ACL can provide. If you simply need to selectively permit some routes based on their prefix or mask, we recommend that you use a route map to map to an ACL (or equivalent prefix list).



Note You must use a standard ACL as the match criterion for your route map. Using an extended ACL will not work, and your routes will never be redistributed. We recommend that you number clauses in intervals of 10 to reserve numbering space in case you need to insert clauses in the future.

Permit and Deny Clauses

Route maps can have permit and deny clauses. If the match criteria are met for this route map, and the permit keyword is specified, the route is redistributed as controlled by the set actions. If the match criteria are not met, and the permit keyword is specified, the next route map with the same map tag is tested. If a route passes none of the match criteria for the set of route maps sharing the same name, it is not redistributed by that set. If the match criteria are met for the route map and the deny keyword is specified, the route is not redistributed.

The following rules apply:

- If you use an ACL in a route map using a permit clause, routes that are permitted by the ACL are redistributed.
- If you use an ACL in a route map deny clause, routes that are permitted by the ACL are not redistributed.
- If you use an ACL in a route map permit or deny clause, and the ACL denies a route, then the route map clause match is not found and the next route-map clause is evaluated.

Match and Set Clause Values

Each entry in a route map statement contains a combination of match and set clauses. The match clause defines the criteria for whether appropriate packets meet the particular policy (that is, the conditions to be met). The set clause explains how the packets should be routed once they have met the match criteria.

For each route that is being redistributed, the router first evaluates the match criteria of a clause in the route map. If the match criteria succeed, then the route is redistributed or rejected as dictated by the permit or deny clause, and some of its attributes might be modified as defined by the set clause. If the match criteria fail, then this clause is not applicable to the route, and the software proceeds to evaluate the route against the next clause in the route map. Scanning of the route map continues until a clause is found whose match clause matches the route or until the end of the route map is reached.

A match or set value in each clause can be missed or repeated several times, if one of these conditions exists:

- If several Match Clause values are present in a clause, all must succeed for a given route in order for that route to match the clause (in other words, the logical AND algorithm is applied for multiple match commands).
- If a Match Clause value refers to several objects in one command, any of the objects should match (the logical OR algorithm is applied).
- If a Match Clause value is not present, all routes match the clause.
- If a Set Value is not present in a route map permit clause, then the route is redistributed without modification of its current attributes.



Note Do not configure a Set Value in a route map deny clause because the deny clause prohibits route redistribution—there is no information to modify.

A route map clause without a Match or Set value performs an action. An empty permit clause allows a redistribution of the remaining routes without modification. An empty deny clause does not allow a redistribution of other routes (this is the default action if a route map is completely scanned, but no explicit match is found).

BGP Match and BGP Set Clauses

In addition to the match and set values described above, BGP provides additional match and set capabilities to route maps.

The following route-map match clauses are supported with BGP:

- match AS path access list
- match community
- match policy list

The following route-map set clauses are supported with BGP:

- set AS path
- set community
- set automatic tag
- set local preference
- set weight
- set origin
- set next hop
- set IP prefix list

Creating and Using Route Map Objects

When configuring a policy that requires that you identify a route map, you can select or create route map objects by clicking the **Select** button next to the Route Map field. To create a new route map from the Route Map Object Selector dialog box, click the **Create** button beneath the route map list. You can also create route map objects from the [Policy Object Manager](#) by selecting **Route Map** from the Object Type Selector and then clicking the **New Object** button. For information on the specific fields available when creating a route map object, see [Add or Edit Route Map Object Dialog Boxes](#), on page 148.

Note about the use of Route Map Objects in BGP Policies

Some of the match and set criteria used in route maps are not supported in all BGP subcommands. For example:

The following route map match criteria:

- Match Clause tab > Match first hop interface of route, Match Next Hop (IPv4 and IPv6), Match Route Source (IPv4 and IPv6), Match Metric Route Value, and Match Tag
- BGP Match Clause tab > Match AS path access lists

and the following route map set criteria:

- Set Clause tab > Metric Values (all fields) and Metric Type
- BGP Set Clause tab > Set AS path, Prepend AS path, and Prepend last AS to the AS path

are not supported in the following places:

- BGP policy > IPv4 Address Family:
 - Aggregate Address tab > Attribute Map, Advertise Map, and Suppress Map
 - Neighbor tab > Filtering tab
 - Route Injection tab > Inject Map and Exist Map

Security Manager allows you to use route maps in your BGP configuration even if the route map contains unsupported match or set criteria and you will not receive a warning or error during validation. In such cases, deployment will fail and you will receive an error from the device in the following format: ...%"My-Route-map" used as BGP inbound route-map, nexthop match not supported... .

Please refer to the ASA documentation for guidelines on the match/set criteria supported in route maps used in BGP configuration.

Related Topics

- [Add or Edit Route Map Object Dialog Boxes , on page 148](#)
- [Add or Edit Route Map Entry Dialog Box , on page 149](#)
- [Selecting Objects for Policies](#)
- [Creating Policy Objects](#)
- [Editing Objects](#)
- [Using Category Objects](#)
- [Managing Object Overrides](#)
- [Allowing a Policy Object to Be Overridden](#)

Add or Edit Route Map Object Dialog Boxes

Use the Add/Edit Route Map Object dialog box to create, copy and edit route map policy objects. You can use route maps to define the conditions for redistributing routes from one routing protocol into another, or to enable policy routing.

Navigation Path

Select **Manage > Policy Objects**, then select **Route Map** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- [Understanding Route Map Objects , on page 145](#)
- [Add or Edit Route Map Entry Dialog Box , on page 149](#)
- [Policy Object Manager](#)
- [Selecting Objects for Policies](#)
- [Creating Policy Objects](#)

- [Editing Objects](#)
- [Using Category Objects](#)
- [Managing Object Overrides](#)
- [Allowing a Policy Object to Be Overridden](#)

Field Reference

Table 92: Add/Edit Route Map Object Dialog Box

Element	Description
Name	Enter a meaningful name for the route map object. The route map object name cannot be more than 58 characters. Caution Security Manager allows you to rename these objects even though you cannot rename them on the device. When you rename these objects in Security Manager, the name change is accomplished by negating the existing CLI, and then issuing new CLI to create and assign the object using the new name. This initial negation may cause routing/network issues in your environment. Security Manager will not provide a warning message about these consequences when you rename the object.
Description	An optional description of the object.
Route Map table	The Route Map entries that are defined in the object. <ul style="list-style-type: none"> • To add a Route Map entry, click the Add button to open the Add or Edit Route Map Entry Dialog Box, on page 149. • To edit a Route Map entry, select it and click the Edit button. • To delete a Route Map entry, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects .
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden and Understanding Policy Object Overrides for Individual Devices . If you allow device overrides, you can click the Edit button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.

Add or Edit Route Map Entry Dialog Box

Use the Add/Edit Route Map Entry dialog box to create a new route map entry for a Route Map object or to edit an existing one.

Navigation Path

From the [Add or Edit Route Map Object Dialog Boxes](#), on page 148, click the **Add** button beneath the Route Map table or select an entry in the table and click the **Edit** button.

Related Topics

- [Understanding Route Map Objects](#), on page 145
- [Add or Edit Route Map Entry Dialog Box](#), on page 149
- [Policy Object Manager](#)
- [Selecting Objects for Policies](#)
- [Creating Policy Objects](#)
- [Editing Objects](#)
- [Using Category Objects](#)
- [Managing Object Overrides](#)
- [Allowing a Policy Object to Be Overridden](#)

Field Reference

Table 93: Add/Edit Route Map Entry Dialog Box

Element	Description
Sequence Number	<p>A number, between 0 and 65535, that indicates the position a new route map entry will have in the list of route maps entries already configured for this route map object.</p> <p>Tip We recommend that you number clauses in intervals of at least 10 to reserve numbering space in case you need to insert clauses in the future.</p>
Redistribution	<p>Whether to redistribute a route or not. To allow redistribution for route matches, click Permit. To reject route matches from redistribution, select Deny.</p> <p>If you use an ACL in a route map Permit clause, routes that are permitted by the ACL are redistributed. If you use an ACL in a route map Deny clause, routes that are permitted by the ACL are not redistributed. In addition, if you use an ACL in a route map Permit or Deny clause, and the ACL denies a route, then the route map clause match is not found and the next route map clause is evaluated.</p>
<p>Match Clause Tab</p> <p>Select the Match Clause tab to choose routes to which this clause should be applied, and set the following parameters:</p>	

Element	Description
Match first hop interface of route	<p>Enable or disable matching routes that have their next hop out one of the interfaces specified. Enter or select the interfaces to match. Separate multiple entries with a comma. If you specify more than one interface, then the route can match either interface.</p> <p>Use the ellipsis to open the Interfaces Selector from which you can select one or more interfaces. You can also create new interface roles from the Interfaces Selector. For more information, see Understanding Interface Role Objects.</p>
IPv4	
Match Address	<p>Enable or disable matching of any routes that have a route address or match packet that is passed by one of the access lists specified.</p> <p>For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List Object Dialog Box , on page 159.</p>
Match Next Hop	<p>Enable or disable matching of the next hop address of a route.</p> <p>For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List Object Dialog Box , on page 159.</p>
Match Route Source	<p>Enable or disable matching of the advertising source address of the route.</p> <p>For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List Object Dialog Box , on page 159.</p>
IPv6	

Element	Description
Match Address	<p>Enable or disable matching of any routes that have a route address or match packet that is passed by one of the access lists specified.</p> <p>For IPv6 addresses, choose whether to use an access list or IPv6 Prefix list for matching from the drop-down list and then enter or select the ACL objects or IPv6 Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or IPv6 Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List IPv6 Object Dialog Box, on page 161.</p>
Match Next Hop	<p>Enable or disable matching of the next hop address of a route.</p> <p>For IPv6 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or IPv6 Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or IPv6 Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List IPv6 Object Dialog Box, on page 161.</p>
Match Route Source	<p>Enable or disable matching of the advertising source address of the route.</p> <p>For IPv6 addresses, choose whether to use an access list or IPv6 Prefix list for matching from the drop-down list and then enter or select the ACL objects or IPv6 Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or IPv6 Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List IPv6 Object Dialog Box, on page 161.</p>
Match Metric Route Value	<p>Enable or disable matching the metric of a route. Type the metric values to use for matching in the Match Metric Route Value field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.</p>
Match Tag	<p>Enable or disable matching the security group tag of a route. Type the tag values to use for matching in the Match Tag field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.</p>
Match Route Type	<p>Enable or disable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. When enabled, you can choose more than one route type from the list.</p>

Element	Description
Set Clause Tab	
Select the Set Clause tab to modify the following information, which will be redistributed to the target protocol:	
Note You can specify just the Bandwidth value, all of the values, or none of the values.	
Bandwidth	Metric value or Bandwidth in Kbits per second; an integer value from 0 to 4294967295.
EIGRP Delay	EIGRP route delay, in tens of microseconds. Valid values range from 1 to 4294967295.
EIGRP Reliability	Likelihood of successful packet transmission for EIGRP expressed as a number from 0 to 255. The value 255 means 100 percent reliability; 0 means no reliability.
EIGRP Effective	Effective EIGRP bandwidth of a route expressed as a number from 1 to 255. The value 255 means 100 percent loading.
EIGRP MTU	Minimum MTU size of a route for EIGRP, in bytes. Valid values range from 1 to 4294967295.
Set Metric Type	Select to specify the type of metric for the destination routing protocol, and choose the metric type from the drop-down list: internal, type-1, or type-2.
BGP Match Clause Tab	
Match AS path access lists	Select to enable matching the BGP autonomous system path access list with the specified path access list. If you specify more than one path access list, then the route can match either path access list. Use the ellipsis to open the AS Path Object Selector from which you can select one or more AS path objects. You can also create new AS path objects from the AS Path Object Selector. For more information, see Add or Edit As Path Object Dialog Boxes , on page 163.
Match community	Select to enable matching the BGP community with the specified community. If you specify more than one community, then the route can match either community. Any route that does not match at least one Match community will not be advertised for outbound route maps. Use the ellipsis to open the Community List Object Selector from which you can select one or more Community List objects. You can also create new Community List objects from the Community List Object Selector. For more information, see Add or Edit Community List Object Dialog Box , on page 165. To enable matching the BGP community exactly with the specified community, check the Match the specified community exactly check box.

Element	Description
Match policy list	<p>Select to configure a route map to evaluate and process a BGP policy. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only.</p> <p>Use the ellipsis to open the Policy List Object Selector from which you can select one or more Policy List objects. You can also create new Policy List objects from the Policy List Object Selector. For more information, see Add or Edit Policy List Object Dialog Box , on page 156.</p>
BGP Set Clause Tab	
Select the BGP Set Clause tab to modify the following information, which will be redistributed to the BGP protocol:	
Set AS path	<p>Select to modify an autonomous system path for BGP routes.</p> <ul style="list-style-type: none"> • Select Prepend AS path to prepend an arbitrary autonomous system path string to BGP routes. Usually the local AS number is prepended multiple times, increasing the autonomous system path length. If you specify more than one AS path number then the route can prepend either AS number. • Select Prepend last AS to the AS path to prepend the AS path with the last AS number. Enter a value for the AS number from 1 to 10. • Select Convert route tag into AS path to convert the tag of a route into an autonomous system path.
Set community	<p>Select to set the BGP communities attributes.</p> <ul style="list-style-type: none"> • Select None to remove the community attribute from the prefixes that pass the route map. • Select Specify community to enter a community number, if applicable. Valid values are from 1 to 4294967295. <p>Select Add to the existing communities to add the community to the already existing communities.</p> <ul style="list-style-type: none"> • Select Internet, no-advertise or no-export to use one of the well-known communities.
Set Automatic-tag	Select to automatically compute the tag value.
Set local preference	Select to specify a preference value for the autonomous system path. Enter a value between 0 and 4294967295.
Set weight	Select to specify the BGP weight for the routing table. Enter a value between 0 and 65535.
Set origin	Select to specify the BGP origin code. Valid values are Local IGP and Incomplete.
Next hop IPv4	

Element	Description
Set next hop	Select to specify the output address of packets that fulfill the match clause of a route map: <ul style="list-style-type: none"> • Select Specify IPv4 address to enter the IPv4 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv4 address then the packets can output at either IP address. • Select Use peer address to set the next hop to be the BGP peer address.
Next hop IPv6	
Set next hop	Select to specify the output address of packets that fulfill the match clause of a route map: <ul style="list-style-type: none"> • Select Specify IPv6 address to enter the IPv6 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv6 address then the packets can output at either IP address. You can enter multiple values separated by commas. • Select Use peer address to set the next hop to be the BGP peer address.
Prefix List	
Set IPv4 prefix list	Select to set an IPv4 prefix list. Use the ellipsis to open the Prefix List Object Selector from which you can select one or more Prefix List objects. You can also create new Prefix List objects from the Prefix List Object Selector. For more information, see Add or Edit Prefix List Object Dialog Box , on page 159.
Set IPv6 prefix list	Select to set an IPv6 prefix list. Use the ellipsis to open the Prefix List Object IPv6 Selector from which you can select one or more IPv6 Prefix List objects. You can also create new IPv6 Prefix List objects from the Prefix List Object Selector. For more information, see Add or Edit Prefix List IPv6 Object Dialog Box , on page 161.
Policy Based Routing (PBR) Tab	
Click the Policy Based Routing tab to define policy for traffic flows, and lessening reliance on routes derived from routing protocols. PBR gives you more control over routing by extending and complementing the existing mechanisms provided by routing protocols. PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link.	
Set Default Next-Hop IPv4 Address	Check the Set default next-hop IPv4 address check box to indicate where to output packets that pass a match clause of a route map for policy routing. In the IPv4 Address enter the destination address.
Set Default Next-Hop IPv6 Address	Check the Set default next-hop IPv6 address check box to indicate where to output packets that pass a match clause of a route map for policy routing. In the IPv6 Address enter the destination address.

Element	Description
Recursively find and set Next-Hop IPv4 Address	Check the Recursively find and set next-hop IP address check box and specify an IP address in the IPv4 Address field. In this case, the next-hop IP address need not be on a directly connected subnet.
Set Interfaces	Check the Set interfaces check box and select a destination interface from the Interfaces Selector dialog box.
Set Null0 Interfaces as Default Interface	Check the Set null0 interface as the default interface check box, if there is a need to completely black hole or drop some traffic.
Set do-not-fragment bit to either 0 or 1	Check the Set do-not-fragment bit to either 1 or 0 and then select the appropriate radio button.
Set Differential Service Code point (DSCP) value in QoS bits for IPv4	Check the Set differential service code point (DSCP) value in QoS bits for IPv4 packets check box and either enter a value between 0 and 63 or select a value from the Select Value drop-down list.
Set Differential Service Code point (DSCP) value in QoS bits for IPv6	Check the Set differential service code point (DSCP) value in QoS bits for IPv6 packets check box and either enter a value between 0 and 63 or select a value from the Select Value drop-down list.
Set Adaptive Interface	<p>Check the Set adaptive Interface box, select the metrics from the drop-down list, and specify the interface in the Set Adaptive Interface dialog box to determine the best path for routing PBR traffic. Following are the adaptive interface metrics:</p> <ul style="list-style-type: none"> • cost—Traffic is forwarded based on the priority of the interfaces. • jitter—Traffic is forwarded to the interface that has the lowest jitter value. • lost—Traffic is forwarded to the interface that has the the minimal packet loss. • mos—Traffic is forwarded to the interface that has the maximum mean opinion score (MOS) • rtt—Traffic is forwarded to the interface that has the the minimal round trip time (RTT)

Add or Edit Policy List Object Dialog Box

Use the Add/Edit Policy List Object dialog box to create, copy, and edit policy list policy objects. You can create policy list objects to use when you are configuring route maps (see [Understanding Route Map Objects](#), on page 145).

When a policy list is referenced within a route map, all of the match statements within the policy list are evaluated and processed. Two or more policy lists can be configured with a route map. A policy list can also coexist with any other preexisting match and set statements that are configured within the same route map but outside of the policy list. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only.

Navigation Path

Select **Manage > Policy Objects**, then select **Policy List** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- [Understanding Route Map Objects](#) , on page 145
- [Policy Object Manager](#)
- [Selecting Objects for Policies](#)
- [Creating Policy Objects](#)
- [Editing Objects](#)
- [Using Category Objects](#)
- [Managing Object Overrides](#)
- [Allowing a Policy Object to Be Overridden](#)

Field Reference

Table 94: Add/Edit Policy List Object Dialog Box

Element	Description
Name	<p>The name of the object. Object names are not case-sensitive. For more information, see Creating Policy Objects.</p> <p>Caution Security Manager allows you to rename these objects even though you cannot rename them on the device. When you rename these objects in Security Manager, the name change is accomplished by negating the existing CLI, and then issuing new CLI to create and assign the object using the new name. This initial negation may cause routing/network issues in your environment. Security Manager will not provide a warning message about these consequences when you rename the object.</p>
Description	An optional description of the object.
Basic Tab	
Action	<p>Whether to permit access for matching conditions or not.</p> <p>Note The Action for a policy list object cannot be changed after the initial creation of the policy list object.</p>
Match Interface	<p>Select to distribute routes that have their next hop out of one of the interfaces specified. Enter or select the interfaces to match. Separate multiple entries with a comma. If you specify more than one interface, then the route can match either interface.</p> <p>Use the ellipsis to open the Interfaces Selector from which you can select one or more interfaces. You can also create new interface roles from the Interfaces Selector. For more information, see Understanding Interface Role Objects.</p>

Element	Description
Match Address	<p>Select to redistribute any routes that have a destination address that is permitted by a standard access list or prefix list. Choose whether to use an Access List or Prefix List for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List Object Dialog Box , on page 159.</p>
Match Next-Hop	<p>Select to redistribute any routes that have a next hop router address passed by one of the access lists or prefix lists specified. Choose whether to use an Access List or Prefix List for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List Object Dialog Box , on page 159.</p>
Match Route Source	<p>Select to redistribute routes that have been advertised by routers and access servers at the address specified by the access lists or prefix list. Choose whether to use an Access List or Prefix List for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching.</p> <p>Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes or Add or Edit Prefix List Object Dialog Box , on page 159.</p>
Advanced Tab	
Match AS Path	<p>Select to match a BGP autonomous system path. If you specify more than one AS path, then the route can match either AS path.</p> <p>Use the ellipsis to open the AS Path Object Selector from which you can select one or more AS path objects. You can also create new AS path objects from the AS Path Object Selector. For more information, see Add or Edit As Path Object Dialog Boxes , on page 163.</p>
Match Community Rules	<p>Select to enable matching the BGP community with the specified community. If you specify more than one community, then the route can match either community.</p> <p>Use the ellipsis to open the Community List Object Selector from which you can select one or more Community List objects. You can also create new Community List objects from the Community List Object Selector. For more information, see Add or Edit Community List Object Dialog Box , on page 165.</p> <p>To enable matching the BGP community exactly with the specified community, check the exact-match check box.</p>

Element	Description
Match Metric	Enable or disable matching the metric of a route. Type the metric values to use for matching in the Match Metric field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.
Match Tag	Enable or disable matching the security group tag of a route. Type the tag values to use for matching in the Match Tag field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects .
Allow Value Override per Device	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden and Understanding Policy Object Overrides for Individual Devices .
Overrides Edit button	If you allow device overrides, you can click the Edit button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.

Add or Edit Prefix List Object Dialog Box

Use the Add/Edit Prefix List Object dialog box to create, copy and edit prefix list policy objects. You can create prefix list objects to use when you are configuring route maps (see [Understanding Route Map Objects](#), on page 145), policy maps (see [Add or Edit Policy List Object Dialog Box](#), on page 156), OSPF Filtering (see [Add/Edit Filtering Dialog Box](#), on page 100), or BGP Neighbor Filtering (see [Add/Edit Neighbor Dialog Box](#), on page 12).

Area Border Router (ABR) type 3 link-state advertisement (LSA) filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number.

Navigation Path

Select **Manage > Policy Objects**, then select **Prefix List** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- [Add or Edit Prefix List Entry Dialog Box](#), on page 161
- [Understanding Route Map Objects](#), on page 145
- [Add or Edit Policy List Object Dialog Box](#), on page 156

- [Policy Object Manager](#)
- [Selecting Objects for Policies](#)
- [Creating Policy Objects](#)
- [Editing Objects](#)
- [Using Category Objects](#)
- [Managing Object Overrides](#)
- [Allowing a Policy Object to Be Overridden](#)

Field Reference

Table 95: Add/Edit Prefix List Object Dialog Box

Element	Description
Name	<p>The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects.</p> <p>Caution Security Manager allows you to rename these objects even though you cannot rename them on the device. When you rename these objects in Security Manager, the name change is accomplished by negating the existing CLI, and then issuing new CLI to create and assign the object using the new name. This initial negation may cause routing/network issues in your environment. Security Manager will not provide a warning message about these consequences when you rename the object.</p>
Description	An optional description of the object.
Prefix List table	<p>The prefix list entries that are defined in the object.</p> <ul style="list-style-type: none"> • To add a prefix list entry, click the Add button to open the Add or Edit Prefix List Entry Dialog Box, on page 161. • To edit a prefix list entry, select it and click the Edit button. • To delete a prefix list entry, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects .
Allow Value Override per Device Overrides Edit button	<p>Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden and Understanding Policy Object Overrides for Individual Devices.</p> <p>If you allow device overrides, you can click the Edit button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.</p>

Add or Edit Prefix List Entry Dialog Box

Use the Add/Edit Prefix List Entry dialog box to create a new prefix list entry or edit an existing one.

Navigation Path

From the [Add or Edit Prefix List Object Dialog Box](#), on page 159, click the **Add** button beneath the Prefix List table or select an entry in the table and click the **Edit** button.

Field Reference

Table 96: Add/Edit Prefix List Entry Dialog Box

Element	Description
Action	Select the Permit or Deny radio button to indicate the redistribution access.
Sequence No	(Optional) Unique number that indicates the position a new prefix list entry will have in the list of prefix list entries already configured for this object. If left blank, the sequence number will default to five more than the largest sequence number currently in use.
IP Address	Specify the prefix number in the format of IP address/mask length.
Minimum Prefix Length	(Optional) Enter the minimum prefix length. The value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.
Maximum Prefix Length	(Optional) Enter the maximum prefix length. The value must be greater than or equal to the Minimum Prefix Length, if present, or greater than the mask length if the Minimum Prefix Length is not specified.

Add or Edit Prefix List IPv6 Object Dialog Box

Use the Add/Edit Prefix List IPv6 Object dialog box to create, copy and edit IPv6 prefix list policy objects. You can create IPv6 prefix list objects to use when you are configuring route maps (see [Understanding Route Map Objects](#), on page 145), policy maps (see [Add or Edit Policy List Object Dialog Box](#), on page 156), OSPF Filtering (see [Add/Edit Filtering Dialog Box](#), on page 100), or BGP Neighbor Filtering (see [Add/Edit Neighbor Dialog Box](#), on page 25).

Area Border Router (ABR) type 3 link-state advertisement (LSA) filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number.

Navigation Path

Select **Manage > Policy Objects**, then select **Prefix ListIPv6** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- [Add or Edit Prefix List Entry Dialog Box](#) , on page 161
- [Understanding Route Map Objects](#) , on page 145
- [Add or Edit Policy List Object Dialog Box](#) , on page 156
- [Policy Object Manager](#)
- [Selecting Objects for Policies](#)
- [Creating Policy Objects](#)
- [Editing Objects](#)
- [Using Category Objects](#)
- [Managing Object Overrides](#)
- [Allowing a Policy Object to Be Overridden](#)

Field Reference

Table 97: Add/Edit IPv6 Prefix List Object Dialog Box

Element	Description
Name	The IPv6 Prefix List object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects . Caution Security Manager allows you to rename these objects even though you cannot rename them on the device. When you rename these objects in Security Manager, the name change is accomplished by negating the existing CLI, and then issuing new CLI to create and assign the object using the new name. This initial negation may cause routing/network issues in your environment. Security Manager will not provide a warning message about these consequences when you rename the object.
Description	An optional description of the object.
IPv6 Prefix List table	The IPv6 prefix list entries that are defined in the object. <ul style="list-style-type: none"> • To add an IPv6 prefix list entry, click the Add button to open the Add or Edit IPv6 Prefix List Entry Dialog Box , on page 163. • To edit an IPv6 prefix list entry, select it and click the Edit button. • To delete an IPv6 prefix list entry, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects .

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden and Understanding Policy Object Overrides for Individual Devices . If you allow device overrides, you can click the Edit button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.

Add or Edit IPv6 Prefix List Entry Dialog Box

Use the Add/Edit IPv6 Prefix List Entry dialog box to create a new IPv6 prefix list entry or edit an existing one.

Navigation Path

From the [Add or Edit IPv6 Prefix List Entry Dialog Box](#), on page 163, click the **Add** button beneath the Prefix List table or select an entry in the table and click the **Edit** button.

Field Reference

Table 98: Add/Edit Prefix List Entry Dialog Box

Element	Description
Action	Select the Permit or Deny radio button to indicate the redistribution access.
Sequence No	(Optional) Unique number that indicates the position a new IPv6 prefix list entry will have in the list of IPv6 prefix list entries already configured for this object. If left blank, the sequence number will default to five more than the largest sequence number currently in use. Note Sequence Number must be in the range of 1 to 4294967295
IPv6 Address	Specify the prefix number in the format: IPv6 address/mask length where mask length is less than or equal to 128.
Minimum Prefix Length	(Optional) Enter the minimum prefix length in the range of 1 to 128. The value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.
Maximum Prefix Length	(Optional) Enter the maximum prefix length in the range of 1 to 128. The value must be greater than or equal to the Minimum Prefix Length, if present, or greater than the mask length if the Minimum Prefix Length is not specified.

Add or Edit As Path Object Dialog Boxes

Use the Add/Edit As Path Object dialog box to create, copy and edit autonomous system (AS) path policy objects. You can create AS path objects to use when you are configuring route maps (see [Understanding](#)

[Route Map Objects , on page 145](#)), policy maps (see [Add or Edit Policy List Object Dialog Box , on page 156](#)), or BGP Neighbor Filtering (see [Add/Edit Neighbor Dialog Box , on page 12](#)).

An AS path filter allows you to filter the routing update message by using access lists and look at the individual prefixes within an update message. If a prefix within the update message matches the filter criteria then that individual prefix is filtered out or accepted depending on what action the filter entry has been configured to carry out.



Note AS path object names must be a unique integer from 1-500. If an AS path object is discovered from a device or configuration file that uses the same name as an existing AS path object, the AS path object on Security Manager will be overwritten regardless of the Allow Device Override for Discovered Policy Objects setting on the Security Manager Administration - Discovery page.

Navigation Path

Select **Manage > Policy Objects**, then select **As Path** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- [Add or Edit As Path Entry Dialog Box , on page 165](#)
- [Understanding Route Map Objects , on page 145](#)
- [Add or Edit Policy List Object Dialog Box , on page 156](#)
- [Policy Object Manager](#)
- [Selecting Objects for Policies](#)
- [Creating Policy Objects](#)
- [Editing Objects](#)
- [Using Category Objects](#)
- [Managing Object Overrides](#)
- [Allowing a Policy Object to Be Overridden](#)

Field Reference

Table 99: Add/Edit As Path Object Dialog Box

Element	Description
Name	Enter a name for the AS Path Filter. Specify a unique value between 1 and 500.
Description	An optional description of the object.

Element	Description
AS Path table	<p>The AS path entries that are defined in the object.</p> <ul style="list-style-type: none"> To add an AS path entry, click the Add button to open the Add or Edit As Path Entry Dialog Box , on page 165. To edit an AS path entry, select it and click the Edit button. To delete an AS path entry, select it and click the Delete button. To rearrange the entries, select an entry and then click the Move Up or Move Down button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects .
Allow Value Override per Device Overrides Edit button	<p>Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden and Understanding Policy Object Overrides for Individual Devices.</p> <p>If you allow device overrides, you can click the Edit button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.</p>

Add or Edit As Path Entry Dialog Box

Use the Add/Edit As Path Entry dialog box to create a new autonomous system (AS) path entry or edit an existing one.

Navigation Path

From the [Add or Edit As Path Object Dialog Boxes](#) , on page 163, click the **Add Row** button beneath the As Path table or select an entry and click the **Edit Row** button.

Field Reference

Table 100: Add/Edit As Path Entry Dialog Box

Element	Description
Action	Select the Permit or Deny radio button to indicate the redistribution access.
Reg Exp	Specify the regular expression that defines the AS path filter. For information on the metacharacters you can use to build regular expressions, see Metacharacters Used to Build Regular Expressions .

Add or Edit Community List Object Dialog Box

Use the Add/Edit Community List Object dialog box to create, copy and edit community list policy objects. You can create community list objects to use when you are configuring route maps (see [Understanding Route Map Objects](#) , on page 145) or policy maps (see [Add or Edit Policy List Object Dialog Box](#) , on page 156).

A community is a group of destinations that share some common attribute. You can use community lists to create groups of communities to use in a match clause of a route map. Just like an access list, a series of community lists can be created. Statements are checked until a match is found. As soon as one statement is satisfied, the test is concluded.

Navigation Path

Select **Manage > Policy Objects**, then select **Community List** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- [Add or Edit Community List Entry Dialog Box](#) , on page 167
- [Understanding Route Map Objects](#) , on page 145
- [Add or Edit Policy List Object Dialog Box](#) , on page 156
- [Policy Object Manager](#)
- [Selecting Objects for Policies](#)
- [Creating Policy Objects](#)
- [Editing Objects](#)
- [Using Category Objects](#)
- [Managing Object Overrides](#)
- [Allowing a Policy Object to Be Overridden](#)

Field Reference

Table 101: Add/Edit Community List Object Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects .
Description	An optional description of the object.
Community List table	The community list entries that are defined in the object. <ul style="list-style-type: none"> • To add a community list entry, click the Add button to open the Add or Edit Community List Entry Dialog Box , on page 167. • To edit a community list entry, select it and click the Edit button. • To delete a community list entry, select it and click the Delete button. • To rearrange the entries, select an entry and then click the Move Up or Move Down button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects .

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden and Understanding Policy Object Overrides for Individual Devices . If you allow device overrides, you can click the Edit button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.

Add or Edit Community List Entry Dialog Box

Use the Add/Edit Community List Entry dialog box to create a new community list entry or edit an existing one.

Navigation Path

From the [Add or Edit Community List Object Dialog Box](#), on page 165, click the **Add** button beneath the Community List table or select an entry in the table and click the **Edit** button.

Field Reference

Table 102: Add/Edit Community List Entry Dialog Box

Element	Description
Type	Select the Standard or Expanded radio button to indicate the community rule type. Note You cannot have entries using Standard and entries using Expanded community rule types in the same Community List object.
Action	Select the Permit or Deny radio button to indicate the redistribution access.
Communities	Specify a community number. Valid values can be from 1 to 4294967295 or from 0:1 to 65534:65535.
internet	Select to specify the Internet well-known community. Routes with this community are advertised to all peers (internal and external).
no-advertise	Select to specify the no-advertise well-known community. Routes with this community are not advertised to any peer (internal or external).
no-export	Select to specify the no-export well-known community. Routes with this community are advertised to only peers in the same autonomous system or to only other sub-autonomous systems within a confederation. These routes are not advertised to external peers.
Expressions	For an expanded community list, specify the regular expression. For information on the metacharacters you can use to build regular expressions, see Metacharacters Used to Build Regular Expressions .

