



CHAPTER 14

Configuring Dynamic And Static Routing

The Routing area lets you edit a static route to ensure that the security appliance correctly forwards network packets destined to the host or network. You can also use a static route to override any dynamic routes that are discovered for this host or network by specifying a static route with a lower metric than the discovered dynamic routes. To create a static route for a host or network, you must define the IP address and metric for the hop gateway to which the security appliance will forward packets destined to the selected host or network. You can also define multiple static routes for a host or network.

This section contains the following topics:

- [Dynamic Routing, page 14-1](#)
- [Static Routes, page 14-28](#)
- [ASR Group, page 14-32](#)
- [Proxy ARPs, page 14-33](#)

Dynamic Routing

The Dynamic Routing area contains the following topics:

- [OSPF](#)
- [RIP](#)

OSPF

OSPF is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. OSPF propagates link-state advertisements 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 when 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, if OSPF is operating on public and private areas, and if address filtering is required, then 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 get flooded to the entire AS including public areas.

If NAT is employed but OSPF is only running in public areas, then 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.

You can have two OSPF routing processes and one RIP routing process running on the security appliance at the same time.

For more information about enabling and configuring OSPF, see the following:

- [Setup](#)
- [Interface](#)
- [Static Neighbor](#)
- [Virtual Link](#)
- [Filtering](#)
- [Redistribution](#)
- [Summary Address](#)

Setup

The Setup pane lets you enable OSPF processes, configure OSPF areas and networks, and define OSPF route summarization.

For more information about configuring these areas, see the following:

- [Setup > Process Instances Tab](#)
- [Setup > Area/Networks Tab](#)
- [Setup > Route Summarization Tab](#)

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Setup > Process Instances Tab

You can enable up to two OSPF process instances. Each OSPF process has its own associated areas and networks.

Fields

- OSPF Process 1 and 2 areas—Each area contains the settings for a specific OSPF process.
- Enable this OSPF Process—Check the check box to enable an OSPF process. Uncheck this check box to remove the OSPF process.
- OSPF Process ID—Enter a unique numeric identifier for the OSPF process. This process ID is used internal and does not need to match the OSPF process ID on any other OSPF devices. Valid values are from 1 to 65535.
- Advanced—Opens the [Edit OSPF Process Advanced Properties](#) dialog box, where you can configure the Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Edit OSPF Process Advanced Properties

You can edit process-specific settings, such as the Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings, in the Edit OSPF Process Advanced Properties dialog box.

Fields

- OSPF Process—Displays the OSPF process you are configuring. You cannot change this value.
- Router ID—To use a fixed router ID, enter a router ID in IP address format in the Router ID field. If you leave this value blank, the highest-level IP address on the security appliance is used as the router ID.
- Ignore LSA MOSPF—Check this check box to suppress the sending of system log messages when the security appliance receives Type 6 (MOSPF) LSA packets. This setting is unchecked by default.
- RFC 1583 Compatible—Check this check box to calculate summary route costs per RFC 1583. Uncheck this check box 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 setting is selected by default.
- Adjacency Changes—Contains settings that define the adjacency changes that cause system log messages to be sent.
 - Log Adjacency Changes—Check this check box to cause the security appliance to send a system log message whenever an OSPF neighbor goes up or down. This setting is selected by default.

- Log Adjacency Changes Detail—Check this check box to cause the security appliance to send a system log message whenever any state change occurs, not just when a neighbor goes up or down. This setting is unchecked by default.
- Administrative Route Distances—Contains the settings for the administrative distances of routes based on the route type.
 - Inter Area—Sets the administrative distance for all routes from one area to another. Valid values range from 1 to 255. The default value is 100.
 - Intra Area—Sets the administrative distance for all routes within an area. Valid values range from 1 to 255. The default value is 100.
 - External—Sets 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 100.
- Timers—Contains the settings used to configure LSA pacing and SPF calculation timers.
 - SPF Delay Time—Specifies the time between when OSPF receives a topology change and when the SPF calculation starts. Valid values range from 0 to 65535. The default value is 5.
 - SPF Hold Time—Specifies the hold time between consecutive SPF calculations. Valid values range from 1 to 65534. The default value is 10.
 - LSA Group Pacing—Specifies 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.
- Default Information Originate—Contains the settings used by an ASBR to generate a default external route into an OSPF routing domain.
 - Enable Default Information Originate—Check this check box to enable the generation of the default route into the OSPF routing domain.
 - Always advertise the default route—Check this check box to always advertise the default route. This option is unchecked by default.
 - Metric Value—Specifies the OSPF default metric. Valid values range from 0 to 16777214. The default value is 1.
 - Metric Type—Specifies the external link type associated with the default route advertised into the OSPF routing domain. Valid values are 1 or 2, indicating a Type 1 or a Type 2 external route. The default value is 2.
 - Route Map—(Optional) The name of the route map to apply. The routing process generates the default route if the route map is satisfied.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Setup > Area/Networks Tab

The Area/Networks tab displays the areas, and the networks they contain, for each OSPF process on the security appliance.

Fields

- Area/Networks—Displays information about the areas and the area networks configured for each OSPF process. Double-clicking a row in the table opens the [Add/Edit OSPF Area](#) dialog box for the selected area.
 - OSPF Process—Displays the OSPF process the area applies to.
 - Area ID—Displays the area ID.
 - Area Type—Displays the area type. The area type is one of the following values: Normal, Stub, NSSA.
 - Networks—Displays the area networks.
 - Authentication—Displays the type of authentication set for the area. The authentication type is one of the following values: None, Password, MD5.
 - Options—Displays any options set for the area type.
 - Cost—Displays the default cost for the area.
- Add—Opens the [Add/Edit OSPF Area](#) dialog box. Use this button to add a new area configuration.
- Edit—Opens the [Add/Edit OSPF Area](#) dialog box. Use this button to change the parameters of the selected area.
- Delete—Removes the selected area from the configuration.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit OSPF Area

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

Fields

- OSPF Process—When adding a new area, choose the OSPF process ID for the OSPF process for which the area is being. If there is only one OSPF process enabled on the security appliance, then 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—Contains the settings for the type of area being configured.
 - 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 have the option of preventing summary LSAs (Type 3 and 4) from being flooded into the area by unchecking the Summary check box.
- Summary—When the area being defined is a stub area, unchecking 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 have the option of preventing summary LSAs from being flooded into the area by unchecking the Summary check box. You can also disable route redistribution by unchecking the Redistribute check box and enabling Default Information Originate.
- Redistribute—Uncheck this check box to prevent routes from being imported into the NSSA. This check box is selected by default.
- Summary—When the area being defined is a NSSA, unchecking this check box prevents LSAs from being sent into the stub area. This check box is selected by default for NSSAs.
- Default Information Originate—Check this check box to generate a Type 7 default into the NSSA. This check box is unchecked 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.
- Area Networks—Contains the settings for defining an OSPF area.
 - Enter IP Address and Mask—Contains the settings used to define the networks in the area.
 - IP Address—Enter the IP address 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.
 - Netmask—Choose the network mask for the IP address or host to be added to the area. If adding a host, choose the 255.255.255.255 mask.
 - Add—Adds the network defined in the Enter IP Address and Mask area to the area. The added network appears in the Area Networks table.
 - Delete—Deletes the selected network from the Area Networks table.
 - Area Networks—Displays the networks defined for the area.
 - IP Address—Displays the IP address of the network.
 - Netmask—Displays the network mask for the network.
- Authentication—Contains the settings for OSPF area authentication.
 - 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. The default value is 1.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Setup > Route Summarization Tab

In OSPF, an ABR will advertise networks in one area into another area. If the network numbers in an area are assigned in a way such that they are contiguous, you can configure the ABR to advertise a summary route that covers all the individual networks within the area that fall into the specified range. To define summary address for external routes being redistributed into an OSPF area, see [Summary Address](#).

Fields

- Route Summarization—Displays information about route summaries defined on the security appliance. Double-clicking a row in the table opens the [Add/Edit Route Summarization](#) dialog box for the selected route summary.
 - OSPF Process—Displays the OSPF process ID for the OSPF process associated with the route summary.
 - Area ID—Displays the area associated with the route summary.
 - IP Address—Displays the summary address.
 - Network Mask—Displays the summary mask.
 - Advertise—Displays “yes” when the route summaries are advertised when they match the address/mask pair or “no” when route summaries are suppressed when they match the address/mask pair.
- Add—Opens the [Add/Edit Route Summarization](#) dialog box. Use this button to define a new route summarization.
- Edit—Opens the [Add/Edit Route Summarization](#) dialog box. Use this button to change the parameters of the selected route summarization.
- Delete—Removes the selected route summarization from the configuration.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit Route Summarization

Use the Add Route Summarization dialog box to add a new entry to the Route Summarization table. Use the Edit Route Summarization dialog box to change an existing entry.

Fields

- OSPF Process—Choose the OSPF process the route summary applies to. You cannot change this value when editing an existing route summary entry.
- Area ID—Choose the area ID the route summary applies to. You cannot change this value when editing an existing route summary entry.
- IP Address—Enter the network address for the routes being summarized.
- Network Mask—Choose one of the common network masks from the list or type the mask in the field.
- Advertise—Check this check box to set the address range status to “advertise”. This causes Type 3 summary LSAs to be generated. Uncheck this check box to suppress the Type 3 summary LSA for the specified networks. This check box is checked by default.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Filtering

The Filtering pane displays the ABR Type 3 LSA filters that have been configured 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.

Fields

The Filtering table displays the following information. Double-clicking a table entry opens the [Add/Edit Filtering Entry](#) dialog box for the selected entry.

- OSPF Process—Displays the OSPF process associated with the filter entry.
- Area ID—Displays the ID of the area associated with the filter entry.
- Filtered Network—Displays the network address 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 coming out of an OSPF area.
- Sequence #—Displays 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—Displays the minimum prefix length to be matched.
- Upper Range—Displays the maximum prefix length to be matched.

You can perform the following actions on entries in the Filtering table:

- Add—Opens the [Add/Edit Filtering Entry](#) dialog box for adding a new entry to the Filter table.
- Edit—Opens the [Add/Edit Filtering Entry](#) dialog box for modifying the selected filter.
- Delete—Removes the selected filter from the Filter table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit Filtering Entry

The Add/Edit Filtering Entry dialog box lets you add new filters to the Filter table or to modify an existing filter. Some of the filter information cannot be changed when you edit an existing filter.

Fields

- OSPF Process—Choose the OSPF process associated with the filter entry. If you are editing an existing filter entry, you cannot modify this setting.
- Area ID—Choose the ID of the area associated with the filter entry. If you are editing an existing filter entry, you cannot modify this setting.
- Filtered Network—Enter the address and mask of the network being filtered using CIDR notation (a.b.c.d/m).
- Traffic Direction—Choose the traffic direction being filtered. Choose “Inbound” to filter LSAs coming into an OSPF area or “Outbound” to filter LSAs coming out of an OSPF area. If you are editing an existing filter entry, you cannot modify this setting.
- Sequence #—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—Choose “Permit” to allow the LSA traffic or “Deny” to block the LSA traffic.
- Optional—Contains the optional settings for the filter.
 - 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Interface

The Interface pane lets you configure interface-specific OSPF authentication routing properties. For more information about configuring these properties, see the following:

- [Interface > Authentication Tab](#)
- [Interface > Properties Tab](#)

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Interface > Authentication Tab

The Authentication tab displays the OSPF authentication information for the security appliance interfaces.

Fields

- Authentication Properties—Displays the authentication information for the security appliance interfaces. Double-clicking a row in the table opens the [Edit OSPF Interface Properties](#) dialog box for the selected interface.
 - Interface—Displays the interface name.
 - Authentication Type—Displays the type of OSPF authentication enabled on the interface. The authentication type can be one of the following values:
 - 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.
- Edit—Opens the [Edit OSPF Interface Properties](#) dialog box for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Edit OSPF Interface Authentication

The Edit OSPF Interface Authentication dialog box lets you configure the OSPF authentication type and parameters for the selected interface.

Fields

- **Interface**—Displays the name of the interface for which authentication is being configured. You cannot edit this field.
- **Authentication**—Contains the OSPF authentication options.
 - **None**—Choose this option to disable OSPF authentication.
 - **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).
 - **Area**—(Default) Choose this option to use the authentication type specified for the area (see [Add/Edit OSPF Area](#) for information about configuring area authentication). Area authentication is disabled by default. So, unless you have previously specified an area authentication type, interfaces set to area authentication have authentication disabled until you configure area authentication.
- **Authentication Password**—Contains the settings for entering the password when password authentication is enabled.
 - **Enter Password**—Enter a text string of up to 8 characters.
 - **Re-enter Password**—Reenter the password.
- **MD5 IDs and Keys**—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.
 - **Enter MD5 ID and Key**—Contains the settings for entering MD5 key information.
 - Key ID**—Enter a numerical key identifier. Valid values range from 1 to 255.
 - Key**—An alphanumeric character string of up to 16 bytes.
 - **Add**—Adds the specified MD5 key to the MD5 ID and Key table.
 - **Delete**—Removes the selected MD5 key and ID from the MD5 ID and Key table.
 - **MD5 ID and Key**—Displays the configured MD5 keys and key IDs.
 - Key ID**—Displays the key ID for the selected key.
 - Key**—Displays the key for the selected key ID.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Interface > Properties Tab

The Properties tab displays the OSPF properties defined for each interface in a table format.

Fields

- OSPF Interface Properties—Displays interface-specific OSPF properties. Double-clicking a row in the table opens the [Edit OSPF Interface Properties](#) dialog box for the selected interface.
 - Interface—Displays the name of the interface that the OSPF configuration applies to.
 - Broadcast—Displays “No” if the interface is set to non-broadcast (point-to-point). Displays “Yes” if the interface is set to broadcast. “Yes” is the default setting for Ethernet interfaces.
 - Cost—Displays the cost of sending a packet through the interface.
 - Priority—Displays the OSPF priority assigned to the interface.
 - MTU Ignore—Displays “No” if MTU mismatch detection is enabled. Displays “Yes” if the MTU mismatch detection is disabled.
 - Database Filter—Displays “Yes” if outgoing LSAs are filtered during synchronization and flooding. Displays “No” if filtering is not enabled.
- Edit—Opens the [Edit OSPF Interface Properties](#) dialog box for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Edit OSPF Interface Properties**Fields**

- Interface—Displays the name of the interface for which you are configuring OSPF properties. You cannot edit this field.
- Broadcast—Check this check box to specify that the interface is a broadcast interface. This check box is selected by default for Ethernet interfaces. Uncheck this check box to designate the interface as a point-to-point, non-broadcast interface. Specifying an interface as point-to-point, non-broadcast lets you transmit OSPF routes over VPN tunnels.

When an interface is configured as point-to-point, non-broadcast, the following restrictions apply:

- You can define only one neighbor for the interface.
 - You need to manually configure the neighbor (see [Static Neighbor](#)).
 - You need to define a static route pointing to the crypto endpoint (see [Static Routes](#)).
 - If OSPF over the tunnel is running on the interface, regular OSPF with an upstream router cannot be run on the same interface.
 - You should bind the crypto-map to the interface before specifying the OSPF neighbor to ensure that the OSPF updates are passed through the VPN tunnel. If you bind the crypto-map to the interface after specifying the OSPF neighbor, use the **clear local-host all** command to clear OSPF connections so the OSPF adjacencies can be established over the VPN tunnel.
- Cost—Specify the cost of sending a packet through the interface. The default value is 10.
 - Priority—Specify the OSPF router priority. When two routers connect to a network, both attempt to become the designated router. The devices with the higher router priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router.
Valid values for this setting range from 0 to 255. The default value is 1. Entering 0 for this setting makes the router 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.
 - MTU Ignore—OSPF checks whether neighbors are using the same MTU on a common interface. This check is performed when neighbors exchange DBD packets. If the receiving MTU in the DBD packet is higher than the IP MTU configured on the incoming interface, OSPF adjacency will not be established.
 - Database Filter—Check this check box to filter outgoing LSA interface during synchronization and flooding. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives. In a fully meshed topology, this can waste bandwidth and lead to excessive link and CPU usage. Checking this check box prevents flooding OSPF LSA on the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Edit OSPF Interface Advanced Properties

The Edit OSPF Interface Advanced Properties dialog box lets you change the values for the OSPF hello interval, retransmit interval, transmit delay, and dead interval. Typically, you only need to change these values from the defaults if you are experiencing OSPF problems on your network.

Fields

- Hello Interval—Specifies 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.

- **Retransmit Interval**—Specifies 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 acknowledgement message. If the router receives no acknowledgement, 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. The default value is 5 seconds.
- **Transmit Delay**—Specifies 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.
- **Dead Interval**—Specifies 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Redistribution

The Redistribution pane displays the rules for redistributing routes from one routing domain to another.

Fields

The Redistribution table displays the following information. Double-clicking a table entry opens the [Add/Edit OSPF Redistribution Entry](#) dialog box for the selected entry.

- **OSPF Process**—Displays the OSPF process associated with the route redistribution entry.
- **Protocol**—Displays the source protocol the routes are being redistributed from. Valid entries are the following:
 - **Static**—The route is a static route.
 - **Connected**—The route was established automatically by virtue of having IP enabled on the interface. These routes are redistributed as external to the AS.
 - **OSPF**—The route is an OSPF route from another process.
- **Match**—Displays the conditions used for redistributing routes from one routing protocol to another.
- **Subnets**—Displays “Yes” if subnetted routes are redistributed. Does not display anything if only routes that are not subnetted are redistributed.
- **Metric Value**—Displays 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**—Displays the name of the route map to apply to the redistribution entry.

You can perform the following actions on the Redistribution table entries:

- **Add**—Opens the [Add/Edit OSPF Redistribution Entry](#) dialog box for adding a new redistribution entry.
- **Edit**—Opens the [Add/Edit OSPF Redistribution Entry](#) dialog box for modifying the selected redistribution entry.
- **Delete**—Removes the selected redistribution entry from the Redistribution table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit OSPF Redistribution Entry

The Add/Edit OSPF Redistribution Entry dialog box lets you add a new redistribution rule to or edit an existing redistribution rule in the Redistribution table. Some of the redistribution rule information cannot be changed when you are editing an existing redistribution rule.

Fields

- **OSPF Process**—Choose the OSPF process associated with the route redistribution entry. If you are editing an existing redistribution rule, you cannot change this setting.
- **Protocol**—Choose the source protocol the routes are being redistributed from. You can choose one of the following options:
 - **Static**—The route is a static route.
 - **Connected**—The route was established automatically by virtue of having IP enabled on the interface. Connected routes are redistributed as external to the AS.
 - **OSPF**—The route is an OSPF route from another process.
OSPF—Choose the OSPF process ID for the route being redistributed.
- **Match**—Displays 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:
 - **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.
- Metric Value—Specify 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—Choose “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—Choose this check box to enable the redistribution of subnetted routes. Uncheck this check box to cause only routes that are not subnetted to be redistributed.
- Route Map—Enter the name of the route map to apply to the redistribution entry.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Static Neighbor

The Static Neighbor pane displays manually defined neighbors; it does not display discovered 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 Static Neighbor table.

Fields

- Static Neighbor—Displays information for the static neighbors defined for each OSPF process. Double-clicking a row in the table opens the [Add/Edit OSPF Neighbor Entry](#) dialog box.
 - OSPF Process—Displays the OSPF process associated with the static neighbor.
 - Neighbor—Displays the IP address of the static neighbor.
 - Interface—Displays the interface associated with the static neighbor.
- Add—Opens the [Add/Edit OSPF Neighbor Entry](#) dialog box. Use this button to define a new static neighbor.
- Edit—Opens the [Add/Edit OSPF Neighbor Entry](#) dialog box. Use this button to change the settings for a static neighbor.
- Delete—Removes the selected entry from the Static Neighbor table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit OSPF Neighbor Entry

The Add/Edit OSPF Neighbor Entry dialog box lets you define a new static neighbor or change information for an existing static neighbor.

You must define a static neighbor for each point-to-point, non-broadcast interface.

Restrictions

- You cannot define the same static neighbor for two different OSPF processes.
- You need to define a static route for each static neighbor (see [Static Routes](#)).

Fields

- OSPF Process—Choose the OSPF process associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.
- Neighbor—Enter the IP address of the static neighbor.
- Interface—Choose the interface associated with the static neighbor. If you are editing an existing static neighbor, you cannot change this value.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Summary Address

The Summary Address pane displays information about the summary addresses configured 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.

Fields

The following information appears in the Summary Address table. Double-clicking an entry in the table opens the [Add/Edit OSPF Summary Address Entry](#) dialog box for the selected entry.

- OSPF Process—Displays the OSPF process associated with the summary address.
- IP Address—Displays the IP address of the summary address.
- Netmask—Displays the network mask of the summary address.
- Advertise—Displays “Yes” if the summary routes are advertised. Displays “No” if the summary route is not advertised.
- Tag—Displays 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.

You can perform the following actions on the entries in the Summary Address table:

- Add—Opens the [Add/Edit OSPF Summary Address Entry](#) dialog box for adding new summary address entries.
- Edit—Opens the [Add/Edit OSPF Summary Address Entry](#) dialog box for editing the selected entry.
- Delete—Removes the selected summary address entry from the Summary Address table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit OSPF Summary Address Entry

The Add/Edit OSPF Summary Address Entry dialog box lets you add new entries to or modify existing entries in the Summary Address table. Some of the summary address information cannot be changed when editing an existing entry.

Fields

- OSPF Process—Choose the OSPF process associated with the summary address. You cannot change this information when editing an existing entry.
- IP Address—Enter the IP address of the summary address. You cannot change this information when editing an existing entry.
- Netmask—Enter the network mask for the summary address, or choose the network mask from the list of common masks. You cannot change this information when editing an existing entry.
- Advertise—Check this check box to advertise the summary route. Uncheck this check box to suppress routes that fall under the summary address. By default this check box is selected.
- Tag—(Optional) 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.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Virtual Link

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.

Fields

The Virtual Link table displays the following information. Doubling-clicking an entry in the table opens the [Add/Edit Virtual Link](#) dialog box for the selected entry.

- OSPF Process—Displays the OSPF process associated with the virtual link.
- Area ID—Displays the ID of the transit area.
- Peer Router ID—Displays the router ID of the virtual link neighbor.
- Authentication—Displays the type of authentication used by the virtual link:
 - None—No authentication is used.
 - Password—Clear text password authentication is used.
 - MD5—MD5 authentication is used.

You can perform the following actions on the entries in the Virtual Link table:

- Add—Opens the [Add/Edit Virtual Link](#) dialog box for adding a new entry to the Virtual Link table.
- Edit—Opens the [Add/Edit Virtual Link](#) dialog box for the selected entry.
- Delete—Removes the selected entry from the Virtual Link table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit Virtual Link

The Add/Edit Virtual Link dialog box lets you define new virtual links or change the properties of existing virtual links.

Fields

- OSPF Process—Choose the OSPF process associated with the virtual link. If you are editing an existing virtual link, you cannot change this value.
- Area ID—Choose the area shared by the neighbor OSPF devices. The selected area cannot be an NSSA or a Stub area. If you are editing an existing virtual link, you cannot change this value.
- Peer Router ID—Enter the router ID of the virtual link neighbor. If you are editing an existing virtual link, you cannot change this value.
- Advanced—Opens the [Advanced OSPF Virtual Link Properties](#) dialog box. You can configure the OSPF properties for the virtual link in this area. These properties include authentication and packet interval settings.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Advanced OSPF Virtual Link Properties

The Advanced OSPF Virtual Link Properties dialog box lets you configure OSPF authentication and packet intervals.

Fields

- Authentication—Contains the OSPF authentication options.
 - None—Choose this option to disable OSPF authentication.
 - 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).
- Authentication Password—Contains the settings for entering the password when password authentication is enabled.
 - Enter Password—Enter a text string of up to 8 characters.
 - Re-enter Password—Reenter the password.
- MD5 IDs and Keys—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.
 - Enter MD5 ID and Key—Contains the settings for entering MD5 key information.
 - Key ID—Enter a numerical key identifier. Valid values range from 1 to 255.
 - Key—An alphanumeric character string of up to 16 bytes.
 - Add—Adds the specified MD5 key to the MD5 ID and Key table.
 - Delete—Removes the selected MD5 key and ID from the MD5 ID and Key table.
 - MD5 ID and Key—Displays the configured MD5 keys and key IDs.

Key ID—Displays the key ID for the selected key.

Key—Displays the key for the selected key ID.

- Intervals—Contains the settings for modifying packet interval timing.
 - Hello Interval—Specifies 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.
 - Retransmit Interval—Specifies 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 acknowledgement message. If the router receives no acknowledgement, 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. The default value is 5 seconds.
 - Transmit Delay—Specifies 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.
 - Dead Interval—Specifies 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 field is four times the interval set by the Hello Interval field.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

RIP

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection. When RIP is enabled on an interface, the interface exchanges RIP broadcasts with neighboring devices to dynamically learn about and advertise routes.

The security appliance support both RIP version 1 and RIP version 2. RIP 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.

You can have two OSPF routing processes and one RIP routing process running on the security appliance at the same time.

Limitations

RIP has the following limitations:

- The security appliance 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.
- You can only enable a single RIP process on the security appliance.

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.

Global Setup

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

Fields

- Enable RIP Routing—Check this check box to enable RIP routing on the security appliance. When you enable RIP, it is enabled on all interfaces. Checking this check box also enables the other fields on this pane. Uncheck this check box to disable RIP routing on the security appliance.
- Enable Auto-summarization—Clear this check box to disable automatic route summarization. Check this check box to reenable automatic route summarization. RIP Version 1 always uses automatic summarization. You cannot disable automatic summarization for RIP Version 1. If you are using RIP Version 2, you can turn off automatic summarization by unchecking this check box. Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is disabled, subnets are advertised.
- Enable RIP version—Check this check box to specify the version of RIP used by the security appliance. If this check box is cleared, then the security appliance sends RIP Version 1 updates and accepts RIP Version 1 & Version 2 updates. This setting can be overridden on a per-interface basis in the [Interface](#) pane.
 - Version 1—Specifies that the security appliance only sends and receives RIP Version 1 updates. Any version 2 updates received are dropped.
 - Version 2—Specifies that the security appliance only sends and receives RIP Version 2 updates. Any version 1 updates received are dropped.
- Enable default information originate—Check this check box to generate a default route into the RIP routing process. You can configure a route map that must be satisfied before the default route can be generated.

- Route-map—Enter the name of the route map to apply. The routing process generates the default route if the route map is satisfied.
- IP Network to Add—Defines a network for the RIP routing process. The network number specified must not contain any subnet information. There is no limit to the number of network you can add to the security appliance configuration. 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.
 - Add—Click this button to add the specified network to the list of networks.
 - Delete—Click this button to removed the selected network from the list of networks.
- Configure interfaces as passive globally—Check this check box to set all interfaces on the security appliance to passive RIP mode. The security appliance listens for RIP routing broadcasts on all interfaces and uses that information to populate the routing tables but do not broadcast routing updates. To set specific interfaces to passive RIP, use the Passive Interfaces table.
- Passive Interfaces table—Lists the configured interfaces on the security appliance. Check the check box in the Passive column for those interfaces you want to operate in passive mode. The other interfaces will still send and receive RIP broadcasts.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Interface

The Interface pane allows you to configure interface-specific RIP settings, such as the version of RIP the interface sends and receives and the authentication method, if any, used for the RIP broadcasts.

Fields

- Interface table—(*Display only*) Each row displays the interface-specific RIP settings for an interface. Double-clicking a row for that entry opens the [Edit RIP Interface Entry](#) dialog box for that interface.
- Edit—Opens the [Edit RIP Interface Entry](#) dialog box for the interface selected in the Interface table.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Edit RIP Interface Entry

The Edit RIP Interface Entry dialog box allows you to configure the interface-specific RIP settings.

Fields

- **Override Global Send Version**—Check this check box to specify the RIP version sent by the interface. You can select the following options:
 - Version 1
 - Version 2
 - Version 1 & 2

Unchecking this check box restores the global setting.

- **Override Global Receive Version**—Check this check box to specify the RIP version accepted by the interface. If a RIP updated from an unsupported version of RIP is received by the interface, it is dropped. You can select the following options:
 - Version 1
 - Version 2
 - Version 1 & 2

Unchecking this check box restores the global setting.

- **Enable Authentication**—Check this check box to enable RIP authentication. Uncheck this check box to disable RIP broadcast authentication.
 - **Key**—The key used by the authentication method. Can contain up to 16 characters.
 - **Key ID**—The key ID. Valid values are from 0 to 255.
 - **Authentication Mode**—You can select the following authentication modes:
 - MD5—Uses MD5 for RIP message authentication.
 - Text—Uses cleartext for RIP message authentication (not recommended).

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Filter Rules

Filter rules allow you to filter the network received in RIP routing updates or sent in RIP routing updates. Each filter rule consists of one or more network rules.

Fields

- **Filter Rules table**—Displays the configured RIP filter rules.

- Add—Clicking this button opens the [Add/Edit Filter Rule](#) dialog box. The new filter rule is added to the bottom of the list.
- Edit—Clicking this button opens the [Add/Edit Filter Rule](#) dialog box for the selected filter rule.
- Delete—Clicking this button deletes the selected filter rule.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit Filter Rule

Use the Add/Edit Filter Rule pane to create filter rules. You can create filter rules that apply to all interfaces or that apply to a specific interface.

Fields

- Direction—Select one of the following directions for the filter to act upon:
 - In—Filters networks on incoming RIP updates.
 - Out—Filters networks from outgoing RIP updates.
- Interface—You can select a specific interface for the filter rule, or you can select the All Interfaces option to apply the filter to all interfaces.
- Action—(*Display only*) Displays Permit if the specified network is not filtered from incoming or outgoing RIP advertisements. Displays Deny if the specified network is to be filtered from incoming or outgoing RIP advertisements.
- IP Address—(*Display only*) Displays the IP address of the network being filtered.
- Netmask—(*Display only*) Displays the network mask applied to the IP address.
- Insert—Click this button to add a network rule above the selected rule in the list. Clicking this button opens the [Network Rule](#) dialog box.
- Edit—Click this button to edit the selected rule. Clicking this button opens the [Network Rule](#) dialog box.
- Add—Click this button to add a network rule below the selected rule in the list. Clicking this button opens the [Network Rule](#) dialog box.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Network Rule

The Network Rule pane allows you to configure permit and deny rules for specific networks in a filter rule.

Fields

- Action—Select Permit to allow the specified network to be advertised in RIP updates or accepted into the RIP routing process. Select Deny to prevent the specified network from being advertised in RIP updates or accepted into the RIP routing process.
- IP Address—Type IP address of the network being permitted or denied.
- Netmask—Specify the network mask applied to the network IP address. You can type a network mask into this field or select one of the common masks from the list.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Route Redistribution

The Route Redistribution pane displays the routes that are being redistributed from other routing processes into the RIP routing process.

Fields

- Protocol—(*Display only*) Displays the routing protocol being redistributed into the RIP routing process:
 - Static—Static routes.
 - Connected—Directly connected networks.
 - OSPF—Networks discovered by the specified OSPF routing process.
- Metric—The RIP metric being applied to the redistributed routes.
- Match—(*Display only*) Displays the type of OSPF routes being redistributed into the RIP routing process. If the Match column is blank for an OSPF redistribution rule, Internal, External 1, and External 2 routes are redistributed into the RIP routing process.
- Route Map—(*Display only*) Displays the name of the route map, if any, being applied to the redistribution. Route maps are used to specify with greater detail which routes from the specified routing process are redistributed into RIP.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Add/Edit Route Redistribution

Use the Add Route Redistribution dialog box to add a new redistribution rule. Use the Edit Route Redistribution dialog box to change an existing rule.

Fields

- Protocol—Choose the routing protocol to redistribute into the RIP routing process:
 - Static—Static routes.
 - Connected—Directly connected networks.
 - OSPF and OSPF ID—Routes discovered by the OSPF routing process. If you choose OSPF, you must also enter the OSPF process ID. Additionally, you can select the specific types of OSPF routes to redistribute from the Match area.
- Route Map—Specifies the name of a route map that must be satisfied before the route can be redistributed into the RIP routing process.
- Configure Metric Type—Check this checkbox to specify a metric for the redistributed routes. If not specified, the routes are assigned a metric of 0.
 - Transparent—Choose this option
 - Value—Choose this to assign a specific metric value. You can enter a value from 0 to 16.
- Match—If you are redistributing OSPF routes into the RIP routing process, you can choose specific types of OSPF routes to redistribute by checking the check box next to the route type. If you do not check any route types, Internal, External 1, and External 2 routes are redistributed by default.
 - Internal—Routes internal to the 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 an NSSA are redistributed.
 - NSSA External 2—Type 2 routes external to an NSSA are redistributed.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	—	•	—	—

Static Routes

Multiple context mode does not support dynamic routing, so you must define static routes for any networks to which the security appliance is not directly connected.

In transparent firewall mode, for traffic that originates on the security appliance and is destined for a non-directly connected network, you need to configure either a default route or static routes so the security appliance knows out of which interface to send traffic. Traffic that originates on the security appliance might include communications to a syslog server, Websense or N2H2 server, or AAA server. If you have servers that cannot all be reached through a single default route, then you must configure static routes.

The simplest option is to configure a default route to send all traffic to an upstream router, relying on the router to route the traffic for you. However, in some cases the default gateway might not be able to reach the destination network, so you must also configure more specific static routes. For example, if the default gateway is on the outside interface, the default route cannot direct traffic to any inside networks that are not directly connected to the security appliance.

You can also use static route in conjunction with dynamic routing protocols to provide a floating static route that is used when the dynamically discovered route goes down. If you create a static route with an administrative distance greater than the administrative distance of the dynamic routing protocol, then a route to the specified destination discovered by the routing protocol takes precedence over the static route. The static route is used only if the dynamically discovered route is removed from the routing table.

Static routes remain in the routing table even if the specified gateway becomes unavailable (see [Static Route Tracking, page 14-29](#), for the exception to this). If the specified gateway becomes unavailable, you need to remove the static route from the routing table manually. However, static routes are removed from the routing table if the associated interface on the security appliance goes down. They are reinstated when the interface comes back up.

You can define up to three equal cost routes to the same destination per interface. ECMP is not supported across multiple interfaces. With ECMP, the traffic is not necessarily divided evenly between the routes; traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses.

The default route identifies the gateway IP address to which the security appliance sends all IP packets for which it does not have a learned or static route. A default route is simply a static route with 0.0.0.0/0 as the destination IP address. Routes that identify a specific destination take precedence over the default route.

You can define up to three equal cost default route entries per device. Defining more than one equal cost default route entry causes the traffic sent to the default route to be distributed among the specified gateways. When defining more than one default route, you must specify the same interface for each entry.

If you attempt to define more than three equal cost default routes, or if you attempt to define a default route with a different interface than a previously defined default route, you will receive an error message.

You can define a separate default route for tunneled traffic along with the standard default route. When you create a default route with the **tunneled** option, all encrypted traffic that arrives on the security appliance and that cannot be routed using learned or static routes is sent to this route. Otherwise, if the traffic is not encrypted, the standard default route entry is used. You cannot define more than one default route with the **tunneled** option; ECMP for tunneled traffic is not supported.

For more information about viewing and configuring static and default routes with ASDM, see [Field Information for Static Routes, page 14-30](#).

Static Route Tracking

It is not always possible to use dynamic routing protocols on the security appliance, such as when the security appliance is in multiple context mode or transparent mode. In these cases, you must use static routes.

One of the problems with static routes is that there is no inherent mechanism for determining if the route is up or down. They remain in the routing table even if the next hop gateway goes down. They are only removed from the routing table if the associated interface on the security appliance goes down.

The static route tracking feature provides a method for tracking the availability of a static route and installing a backup route if the primary route should fail. This allows you to, for example, define a default route to an ISP gateway and a backup default route to a secondary ISP in case the primary ISP becomes unavailable.

The security appliance does this by associating a static route with a monitoring target that you define. It monitors the target using ICMP echo requests. If an echo reply is not received within a specified time period, the object is considered down and the associated route is removed from the routing table. A previously configured backup route is used in place of the removed route.

When selecting a monitoring target, you need to make sure that it can respond to ICMP echo requests. The target can be any network object that responds to ICMP echo requests. Consider choosing:

- the ISP gateway (for dual ISP support) address
- the next hop gateway address (if you are concerned about the availability of the gateway)
- a server, such as a AAA server, that the security appliance needs to communicate with
- a persistent network object on the destination network (a desktop or notebook computer that may be shut down at night is not a good choice)

For more information about configuring static route tracking, see [Configuring Static Route Tracking, page 14-29](#). To monitor the static route tracking process, see [interface connection, page 39-9](#).

Configuring Static Route Tracking

This procedure provides an overview of configuring static route tracking. For specific information about the various fields used to configure this feature, see [Field Information for Static Routes, page 14-30](#)

To configure tracking for a static route, perform these steps:

-
- Step 1** Choose a target of interest. Make sure the target responds to echo requests.
 - Step 2** Open the Static Routes page. Go to **Configuration > Routing > Static Routes**.
 - Step 3** Click **Add** to configure a static route that is to be used based on the availability of your selected target of interest. You must enter the Interface, IP Address, Mask, Gateway, and Metric for this route. See [Add/Edit Static Route, page 14-31](#), for more information about these fields.
 - Step 4** Choose **Tracked** in the Options area for this route.
 - Step 5** Configure the tracking properties. You must enter a unique Track ID, a unique SLA ID, and the IP address of your target of interest. See [Add/Edit Static Route, page 14-31](#), for more information about these fields.
 - Step 6** (Optional) To configure the monitoring properties, click **Monitoring Options** in the Add Static Route dialog box. See [Route Monitoring Options, page 14-32](#), for more information about the monitoring properties.
 - Step 7** Click **OK** to save your changes.

The monitoring process begins as soon as you save the tracked route.

- Step 8** Create a secondary route. The secondary route is a static route to the same destination as the tracked route, but through a different interface or gateway. You must assign this route a higher administrative distance (metric) than your tracked route.
-

Field Information for Static Routes

For information about a specific pane, see the following topics:

- [Static Routes, page 14-30](#)
- [Add/Edit Static Route, page 14-31](#)
- [Route Monitoring Options, page 14-32](#)

Static Routes

The Static Route pane lets you create static routes that will access networks connected to a router on any interface. To enter a default route, set the IP address and mask to 0.0.0.0, or the shortened form of 0.

If an IP address from one security appliance interface is used as the gateway IP address, the security appliance will ARP the designated IP address in the packet instead of ARPing the gateway IP address.

Leave the Metric at the default of 1 unless you are sure of the number of hops to the gateway router.

Fields

The Static Route pane shows the Static Route table:

- Interface—(*Display only*) Lists the internal or external network interface name enabled in Interfaces.
- IP Address—(*Display only*) Lists the internal or external network IP address. Use **0.0.0.0** to specify a default route. The **0.0.0.0** IP address can be abbreviated as **0**.
- Netmask—(*Display only*) Lists the network mask address that applies to the IP address. Use **0.0.0.0** to specify a default route. The **0.0.0.0** netmask can be abbreviated as **0**.
- Gateway IP—(*Display only*) Lists the IP address of the gateway router, which is the next hop address for this route.
- Metric—(*Display only*) Lists the administrative distance of the route. The default is 1 if a metric is not specified.
- Options—(*Display only*) Displays any options specified for the static route.
 - None—No options are specified for the static route.
 - Tunneled—Specifies route as the default tunnel gateway for VPN traffic. Used only for default route. You can only configure one tunneled route per device. The tunneled option is not supported under transparent mode.
 - Tracked—Specifies that the route is tracked. The tracking object ID and the address of the tracking target are also displayed. The tracked option is only supported in single, routed mode.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—

Add/Edit Static Route

Use the Add/Edit Static Route dialog box to configure the static route properties. This dialog box is available from both the Static Routes screen in the Startup Wizard and the Configuration > Routing > Static Route pane.

Fields

- Interface Name—Select the egress interface for the route.
- IP Address—Specifies the internal or external network IP address. Use **0.0.0.0** to specify a default route. The **0.0.0.0** IP address can be abbreviated as **0**.
- Mask—Specifies the network mask address that applies to the IP address. Use **0.0.0.0** to specify a default route. The **0.0.0.0** netmask can be abbreviated as **0**.
- Gateway IP—Specifies the IP address of the gateway router, which is the next hop address for this router.
- Metric—Lets you specify the administrative distance of the route. The default is **1** if a metric is not specified.

The following options are available for static routes. You can select only one of these options for a static route. By default, no option (None) is selected.

- None—No options are specified for the static route.
- Tunneled—Used only for default route. Only one default tunneled gateway is allowed per security appliance. Tunneled option is not supported under transparent mode.
- Tracked—Select this option to specify that the route is tracked. Specifying this option starts the route tracking process.
 - Track ID—A unique identifier for the route tracking process.
 - Track IP Address/DNS Name—Enter the IP address or hostname of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.
 - SLA ID—A unique identifier for the SLA monitoring process.
 - Monitor Options—Click this button to open the [Route Monitoring Options](#) dialog box. In the [Route Monitoring Options](#) dialog box you can configure the parameters of the tracked object monitoring process.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	•	•	•	—

Route Monitoring Options

Use the Route Monitoring Options dialog box to change the tracking object monitoring properties.

Fields

- **Frequency**—Enter how often, in seconds, the security appliance should test for the presence of the tracking target. The default value is 60 seconds. Valid values are from 1 to 604800 seconds.
- **Threshold**—Enter the amount of time, in milliseconds, that indicates an over-threshold event. This value cannot be more than the timeout value.
- **Timeout**—Enter the amount of time, in milliseconds, the route monitoring operation should wait for a response from the request packets. The default value is 5000 milliseconds. Valid values are from 0 to 604800000 milliseconds.
- **Data Size**—Enter the size of data payload to use in the echo request packets. The default value is 28. Valid values are from 0 to 16384.



Note This setting specifies the size of the payload only; it does not specify the size of the entire packet.

- **ToS**—Enter a value for the type of service byte in the IP header of the echo request. The default value is 0. Valid values are from 0 to 255.
- **Number of Packets**—The number of echo requests to send for each test. The default value is 1. Valid values are from 1 to 100.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
			Multiple	
Routed	Transparent	Single	Context	System
•	—	•	—	—

ASR Group

Use the ASR Group screen to assign asynchronous routing group ID numbers to interfaces.

In some situations, return traffic for a session may be routed through a different interface than it originated from. In failover configurations, return traffic for a connection that originated on one unit may return through the peer unit. This most commonly occurs when two interfaces on a single security

appliance, or two security appliances in a failover pair, are connected to different service providers and the outbound connection does not use a NAT address. By default, the security appliance drops the return traffic because there is no connection information for the traffic.

You can prevent the return traffic from being dropped using an ASR Group on interfaces where this is likely to occur. When an interface configured with an ASR Group receives a packet for which it has no session information, it checks the session information for the other interfaces that are in the same group.

**Note**

You must enable Stateful Failover for session information to be passed from the standby failover group to the active failover group.

If it does not find a match, the packet is dropped. If it finds a match, then one of the following actions occurs:

- If the incoming traffic originated on a peer unit in a failover configuration, some or all of the layer 2 header is rewritten and the packet is redirected to the other unit. This redirection continues as long as the session is active.
- If the incoming traffic originated on a different interface on the same unit, some or all of the layer 2 header is rewritten and the packet is reinjected into the stream.

Fields

The **ASR Group** table displays the following information for each interface on the security appliance:

- **Interface**—Displays the name of the interface on the security appliance.
- **ASR Group ID**—Displays the number of the ASR Group the interface belongs to. If the interface has not been assigned an ASR Group number, this column displays "-- None --". Valid values are from 1 to 32.

To assign an ASR Group number to an interface, click the **ASR Group ID** cell in the row of the desired interface. A list of valid ASR Group number appears. Select the desired ASR Group number from the list. You can assign a maximum of 8 interfaces to a single ASR Group. If other contexts have interfaces assigned to an ASR Group, those interface count against the total of 8, even for the context currently being configured.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	—	•	—

Proxy ARPs

In rare circumstances, you might want to disable proxy ARP for global addresses.

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. 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.”

Proxy ARP is when a device responds to an ARP request with its own MAC address, even though the device does not own the IP address. The security appliance uses proxy ARP when you configure NAT and specify a global address that is on the same network as the security appliance interface. The only way traffic can reach the hosts is if the security appliance uses proxy ARP to claim that the security appliance MAC address is assigned to destination global addresses.

Fields

- Interface—Lists the interface names.
- Proxy ARP Enabled—Shows whether proxy ARP is enabled or disabled for NAT global addresses, Yes or No.
- Enable—Enables proxy ARP for the selected interface. By default, proxy ARP is enabled for all interfaces.
- Disable—Disables proxy ARP for the selected interface.

Modes

The following table shows the modes in which this feature is available:

Firewall Mode		Security Context		
Routed	Transparent	Single	Multiple	
			Context	System
•	•	•	•	—