



# RIP for Firepower Threat Defense

---

This chapter describes how to configure the Firepower Threat Defense to route data, perform authentication, and redistribute routing information, using the Routing Information Protocol (RIP).

- [About RIP, on page 1](#)
- [Requirements and Prerequisites for RIP, on page 3](#)
- [Guidelines for RIP, on page 3](#)
- [Configure RIP, on page 4](#)

## About RIP

The Routing Information Protocol, or RIP, as it is more commonly called, is one of the most enduring of all routing protocols. RIP has four basic components: routing update process, RIP routing metrics, routing stability, and routing timers. Devices that support RIP send routing-update messages at regular intervals and when the network topology changes. These RIP packets include information about the networks that the devices can reach, as well as the number of routers or gateways that a packet must travel through to reach the destination address. RIP generates more traffic than OSPF, but is easier to configure.

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 Firepower Threat Defense device supports 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 Firepower Threat Defense device receives reliable routing information from a trusted source.

RIP has advantages over static routes because the initial configuration is simple, and you do not need to update the configuration when the topology changes. The disadvantage to RIP is that there is more network and processing overhead than in static routing.

## Routing Update Process

RIP sends routing-update messages at regular intervals and when the network topology changes. When a router receives a routing update that includes changes to an entry, it updates its routing table to reflect the new route. The metric value for the path is increased by 1, and the sender is indicated as the next hop. RIP routers maintain only the best route (the route with the lowest metric value) to a destination. After updating

its routing table, the router immediately begins transmitting routing updates to inform other network routers of the change. These updates are sent independently of the regularly scheduled updates that RIP routers send.

## RIP Routing Metric

RIP uses a single routing metric (hop count) to measure the distance between the source and a destination network. Each hop in a path from source to destination is assigned a hop count value, which is typically 1. When a router receives a routing update that contains a new or changed destination network entry, the router adds 1 to the metric value indicated in the update and enters the network in the routing table. The IP address of the sender is used as the next hop.

## RIP Stability Features

RIP prevents routing loops from continuing indefinitely by implementing a limit on the number of hops allowed in a path from the source to a destination. The maximum number of hops in a path is 15. If a router receives a routing update that contains a new or changed entry, and if increasing the metric value by 1 causes the metric to be infinity (that is, 16), the network destination is considered unreachable. The downside of this stability feature is that it limits the maximum diameter of a RIP network to less than 16 hops.

RIP includes a number of other stability features that are common to many routing protocols. These features are designed to provide stability despite potentially rapid changes in network topology. For example, RIP implements the split horizon and hold-down mechanisms to prevent incorrect routing information from being propagated.

## RIP Timers

RIP uses numerous timers to regulate its performance. Following are the timer stages for RIP:

- **Update**—The routing-update timer is the interval between periodic routing updates. This is how often the device sends routing updates. Generally, it is set to 30 seconds, with a small random amount of time added whenever the timer is reset. This is done to help prevent congestion, which could result from all routers simultaneously attempting to update their neighbors.
- **Invalid**—Each routing table entry has a route-timeout timer associated with it. This is the number of seconds since the device received the last valid update. When the route-timeout timer expires, the route is marked invalid but is retained in the table until the route-flush timer expires. Once this timer expires, the route goes into holddown. The default is 180 seconds (3 minutes).
- **Holddown**—The holddown period is the number of seconds the system waits before accepting any new updates for the route that is in holddown (that is, routes that have been marked invalid). The default is 180 seconds (3 minutes).
- **Flush**—The route-flush timer is the number of seconds since the system received the last valid update until the route is discarded and removed from the routing table. The default is 240 seconds (4 minutes).

As an example, when the interface on an adjacent router goes down, the system no longer receives routing updates from the adjacent router. At this time, the Invalid and Flush timers start increasing. In the first 180 seconds, nothing will happen. After 180 seconds, the invalid timer expires, making the route invalid, and the Holddown timer starts and holds the route for another 60 seconds. If there is still no update regarding the interface status on the adjacent router (that is, it is still down), then the route enters into the Flush state where in total the system has waited for 240 seconds from the last update (180 seconds for the Invalid timer and 60 seconds for Holddown timer), and the system flushes the route. Even if the adjacent routers interface comes

up immediately, the system does not accept a routing update until the Holddown timer completes the remaining 120 seconds.

## Requirements and Prerequisites for RIP

### Model Support

FTD  
FTDv

### Supported Domains

Any

### User Roles

Admin  
Network Admin

## Guidelines for RIP

### IPv6 Guidelines

Does not support IPv6.

### Additional Guidelines

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 Firepower Threat Defense device 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.

### Limitations

- The Firepower Threat Defense device 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 Firepower Threat Defense device.

# Configure RIP

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection.

- 
- Step 1** Choose **Devices > Device Management**, and edit the Firepower Threat Defense device.
- Step 2** Select **Routing**.
- Step 3** Select **RIP** from the table of contents.
- Step 4** Select the **Enable RIP** check box to configure the RIP settings.
- Step 5** Select the RIP versions for sending and receiving RIP updates from the **RIP Version** drop-down list.
- Step 6** (Optional) Select the **Generate Default Route** check box to generate a default route for distribution, based on the route map that you specify.
- a) Specify a route map name to use for generating default routes, in the **Route Map** field.  
The default route 0.0.0.0/0 is generated for distribution over a certain interface, when the route map, specified in the **Route Map** field, is present.
- Step 7** When Send and Receive Version 2 is the chosen RIP Version, the **Enable Auto Summary** option is available. When the **Enable Auto Summary** checkbox is 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.
- Step 8** Click **Networks**. Define one or more networks for RIP routing. Enter IP address(es), or enter or select the desired Network/Hosts objects. There is no limit to the number of networks you can add to the security appliance configuration. Any interface that belongs to a network defined by this command, will participate in the RIP routing process. 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.
- Note** RIP only supports IPv4 objects.
- Step 9** (Optional) Click **Passive Interface**. Use this option to specify passive interfaces on the 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.
- Step 10** Click **Redistribution** to manage redistribution routes. These are the routes that are being redistributed from other routing processes into the RIP routing process.
- a) Click **Add** to specify redistribution routes.
  - b) Select the routing protocol to redistribute into the RIP routing process, in the **Protocol** drop-down list.
 

**Note** For the OSPF protocol, specify a process ID. Similarly, specify an AS path for BGP. When you choose the Connected option in the **Protocol** drop-down list, you can redistribute, directly connected networks into the RIP routing process.
  - c) (Optional) If you are redistributing OSPF routes into the RIP routing process, you can select specific types of OSPF routes to redistribute in the **Match** drop-down list. Ctrl-click to select multiple types:
    - 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

**Note** The default is match Internal, External 1, and External 2

- d) Select the RIP metric type to apply to the redistributed routes in the **Metric** drop-down list. The two choices are:
- Transparent – Use the current route metric
  - Specified Value – Assign a specific metric value. Enter a specific value from 0-16, in the **Metric Value** field.
  - None – No metric is specified. Do not use any metric value, to apply to redistributed routes.
- e) (Optional) Enter the name of a route map that must be satisfied, in the **Route Map** field before the route can be redistributed into the RIP routing process. Routes are redistributed only if IP address matches an allow statement in the route map address list.
- f) Click **OK**.

### Step 11

(Optional) Click **Filtering** to manage filters for the RIP policy. In this section, filters are used to prevent routing updates through an interface, control the advertising of routes in routing updates, control the processing of routing updates and filtering sources of routing updates.

- a) Click **Add** to add RIP filters.
- b) Select the type of traffic to be filtered - Inbound or Outbound in the **Traffic Direction** field.

**Note** If traffic direction is inbound, you can only define an Interface filter.

- c) Specify whether the filter is based on an Interface or a Route, by selecting appropriate in the **Filter On** field. 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:
- Static – Only static routes are filtered.
  - Connected – Only connected routes are filtered.
  - OSPF – Only OSPFv2 routes discovered by the specified OSPF process are filtered. Enter the Process ID of the OSPF process to be filtered.
  - BGP – Only BGPv4 routes discovered by the specified BGP process are filtered. Enter the AS path of the BGP process to be filtered.
- d) In the **Access List** field, 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.
- e) Click **OK**.

### Step 12

(Optional) Click **Broadcast** to add or edit interface configurations. Using Broadcastf, you can override the global RIP versions to send or receive per interface. You can also define the authentication parameters per interface if you want to implement authentication to ensure valid RIP updates.

- a) Click **Add** to add interface configurations.
- b) Enter or Select an interface defined on this appliance in the **Interface** field.
- c) In the Send option, select the appropriate boxes to specify sending updates using the RIP **Version 1**, **Version 2**, or both. These options let you override, for the specified interface, the global Send versions specified .

- d) In the Receive option, select the appropriate boxes to specify accepting updates using the RIP **Version 1**, **Version 2**, or both. These options let you override, for the specified interface, the global Receive versions specified .
- e) Select the **Authentication** used on this interface for RIP broadcasts.

- None – No authentication
- MD5 – Employ MD5
- Clear Text – Employ clear-text authentication

If you choose MD5 or Clear Text, you must also provide the following authentication parameters.

- 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

- f) Click **OK**.
-