Chapter 21

RIP Configuration Guidelines

To configure the Routing Information Protocol (RIP), you include the following statements:

```
protocols {
    rip {
        any-sender;
        authentication-key password;
        authentication-type type;
        (check-zero | no-check-zero);
        graceful-restart {
            disable;
            restart-time seconds;
        }
        holddown seconds;
        import [ policy-names ];
        message-size number;
        metric-in metric;
        receive receive-options;
        rib-group group-name;
        route-timeout seconds;
        send send-options;
        update-interval seconds;
        traceoptions {
            file name <replace> <size size> <files number> <no-stamp>
                <(world-readable | no-world-readable)>;
            flag flag <flag-modifier> <disable>;
        }
        group group-name {
            bfd-liveness-detection {
                detection-time {
                    threshold milliseconds;
                }
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                transmit-interval {
                    threshold milliseconds;
                    minimum-interval milliseconds;
                }
                multiplier number;
                version (0 | 1 | automatic);
            }
            export [ policy-names ];
            metric-out metric
```

```
            preference number;
            route-timeout seconds;
            update-interval seconds;
            neighbor neighbor-name {
              authentication-key password;
              authentication-type type;
              bfd-liveness-detection {
                detection-time {
                  threshold milliseconds;
                }
                minimum-interval milliseconds;
                minimum-receive-interval milliseconds;
                transmit-interval {
                  threshold milliseconds;
                  minimum-interval milliseconds;
                }
                multiplier number;
                version (0 | 1 | automatic);
              }
              (check-zero | no-check-zero);
              import [ policy-names ];
              message-size number;
              metric-in metric;
              metric-out metric;
              receive receive-options;
              route-timeout seconds;
              send send-options;
              update-interval seconds;
            }
          }
        }
    }
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

By default, RIP is disabled.

To have a router exchange routes with other routers, you must configure RIP groups and neighbors. RIP routes received from routers not configured as RIP neighbors are ignored. Likewise, RIP routes are advertised only to routers configured as RIP neighbors, with an appropriate RIP export policy applied.

This chapter discusses the following topics:

Minimum RIP Configuration

For a router to accept RIP routes, you must include at least the rip, group, and neighbor statements. Routes received from routers that are not configured as neighbors are ignored. All other RIP configuration statements are optional. This minimum configuration defines one neighbor. Include one neighbor statement for each interface on which you want to receive routes. The local router imports all routes by default from this neighbor and does not advertise routes. The router can receive both version 1 and version 2 update messages, with 25 route entries per message. For routing instances, include the statements at the [edit routing-instances *routing-instance-name* protocols rip] hierarchy level.

```
protocols {
    rip {
        group group-name {
            neighbor interface-name {
            }
        }
    }
}
```

NOTE: When you configure RIP on an interface, you must also configure family inet at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level. For more information about the family inet statement, see the *JUNOS Network Interfaces Configuration Guide*.

Defining RIP Global Properties

To define RIP global properties, which apply to all RIP neighbors, include one or more of the following statements. These statements are explained in separate sections.

```
authentication-key password;
authentication-type type;
(check-zero | no-check-zero);
import [ policy-names ];
message-size number;
metric-in metric;
receive receive-options;
rib-group group-name;
send send-options;
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

Defining RIP Neighbor Properties

To define neighbor-specific properties, include one or more of the following statements. The statements are explained in separate sections.

```
neighbor neighbor-name {
    authentication-key password;
    authentication-type type;
    bfd-liveness-detection {
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
        version (0 | 1 | automatic);
    }
    (check-zero | no-check-zero);
    import [ policy-names ];
    message-size number;
    metric-in metric;
    receive receive-options;
    send send-options;
}
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

Configuring Authentication

You can configure the router to authenticate RIP route queries. By default, authentication is disabled. You can use the following authentication method:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.
- MD5 authentication—Creates an encoded checksum that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet's MD5 checksum.

To enable authentication and specify an authentication method and password, include the authentication-key and authentication-type statements:

```
authentication-key password;
authentication-type type;
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

The password can be up to 16 contiguous characters and can include any ASCII strings.

Modifying the Incoming Metric

By default, RIP imports routes from the neighbors configured with the neighbor statement. These routes include those learned from RIP as well as those learned from other protocols. By default, routes that RIP imports from its neighbors have a metric of 1 added to the current route metric.

To change the default metric to be added to incoming routes, include the metric-in statement:

```
metric-in metric;
```

metric can be a value from 1 through 16.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Configuring RIP Timers

You can configure various timers for RIP.

RIP routes expire when either a route timeout limit is met or a route metric reaches infinity, and the route is no longer valid. However, the expired route is retained in the routing table for a time period so that neighbors can be notified that the route has been dropped. This time period is set by configuring the hold-down timer. Upon expiration of the hold-down timer, the route is removed from the routing table.

To configure the hold-down timer for RIP, include the holddown statement:

```
holddown seconds;
```

seconds can be a value from 10 through 180. The default value is 120 seconds.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You can set a route timeout interval. If a route is not refreshed after being installed into the routing table by the specified time interval, the route is removed from the routing table.

To configure the route timeout for RIP, include the route-timeout statement:

```
route-timeout seconds;
```

seconds can be a value from 60 through 360. The default value is 180 seconds.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You can set an update time interval to periodically send out routes learned by RIP to neighbors.

To configure the update time interval, include the update-interval statement:

```
update-interval seconds;
```

seconds can be a value from 10 through 60. The default value is 30 seconds.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Configuring the Number of Route Entries in an Update Message

By default, RIP includes 25 route entries in each update message. To change the number of route entries in an update message, include the message-size statement:

```
message-size number;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

NOTE: To ensure interoperability with routers from other vendors, do not change the default number of route entries in a RIP update message.

Accepting Packets Whose Reserved Fields Are Nonzero

Some of the reserved fields in RIP version 1 packets must be zero, while in RIP version 2 packets most of these reserved fields can contain nonzero values. By default, RIP discards version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.

If you find that you are receiving RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero, you can configure RIP to receive these packets in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453. To receive packets whose reserved fields are nonzero, include the no-check-zero statement:

```
no-check-zero;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Configuring Update Messages

You can configure whether the RIP update messages conform to RIP version 1 only, to RIP version 2 only, or to both versions. You can also disable the sending or receiving of update messages. To configure the sending and receiving of update messages, include the receive and send statements:

```
receive receive-options;
send send-options;
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

Configuring Routing Table Groups

You can install routes learned through RIP into multiple routing tables by configuring a routing table group. RIP routes are installed into each routing table that belongs to that routing table group. To configure a routing table group for RIP routes, include the rib-group statement:

```
rib-group group-name;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Applying Import Policy

To filter routes being imported by the local router from its neighbors, include the import statement and list the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in order (first to last) and the first matching policy is applied to the route. If no match is found, the local router does not import any routes.

```
import [ policy-names ];
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

For more information about creating policies, see the *JUNOS Policy Framework Configuration Guide*.

Configuring Group-Specific Properties

You can group together neighbors that share the same export policy and export metric defaults. You configure group-specific RIP properties by including the group statement at the [edit protocols rip] hierarchy level. Each group must contain at least one neighbor. You should create a group for every export policy you have. To configure neighbors, see "Defining RIP Global Properties" on page 440.

```
[edit protocols rip]
group group-name {
    bfd-liveness-detection {
        detection-time {
            threshold milliseconds;
        }
        minimum-interval milliseconds;
        minimum-receive-interval milliseconds;
        transmit-interval {
            threshold milliseconds;
            minimum-interval milliseconds;
        }
        multiplier number;
        version (0 | 1 | automatic);
    }
    export [ policy-names ];
    preference number;
    metric-out metric;
    neighbor neighbor-options;
}
```

This section discusses the following tasks:

- Applying Export Policy on page 445
- Controlling Route Preference on page 445
- Modifying the Outgoing Metric on page 445

Applying Export Policy

By default, RIP does not export routes it has learned to its neighbors. To have RIP export routes, apply one or more export policies. To apply export policies and to filter routes being exported from the local router to its neighbors, include the export statement and list the name of the policy to be evaluated:

```
export [ policy-names ];
```

To configure export policy globally for all RIP neighbors, include the export statement.

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You can define one or more export policies. If no routes match the policies, the local router does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the metric-in and metric-out values.

For more information about creating policies, see the *JUNOS Policy Framework Configuration Guide*.

Controlling Route Preference

By default, the JUNOS software assigns a preference of 100 to routes that originate from RIP. When the JUNOS software determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table. (For more information about preferences, see "Route Preferences" on page 6.)

To modify the default RIP preference value, include the preference statement:

```
preference preference;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

preference can be a value from 0 through 4,294,967,295 ($2^{32}-1$).

Modifying the Outgoing Metric

If you have included the export statement, RIP exports routes it has learned to the neighbors configured with the neighbor statement.

The metric associated with a RIP route (unless modified by an export policy) is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with a metric-in value of 2 is advertised with a combined metric of 7 when advertised to RIP neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the metric-out statement. The default value for metric-out is 1.

The metric for a route may be modified with an export policy. That metric is seen when the route is exported to the next hop.

To increase the metric for routes advertised outside a group, include the metric-out statement:

```
metric-out metric;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

Configuring Graceful Restart

Graceful restart is disabled by default. You can globally enable graceful restart for all routing protocols at the [edit routing-options] hierarchy level.

You can configure graceful restart parameters specifically for RIP. To do this, include the graceful-restart statement:

```
graceful-restart {
    restart-time seconds;
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

To disable graceful restart for RIP, specify the disable statement. To configure a time period for the restart to finish, specify the restart-time statement.

Configuring the BFD Protocol

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. BFD failure detection times are smaller than RIP detection times, providing faster reaction times to various kinds of failures in the network. You can adjust these timers to be more or less aggressive.

NOTE: To enable BFD for RIP, both sides of the connection must receive an update message from the peer. By default, RIP does not export any routes. Therefore you must enable update messages to be sent by configuring an export policy for routes before a BFD session is triggered.

To enable failure detection, include the bfd-liveness-detection statement:

```
bfd-liveness-detection {
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
    }
```

```
        multiplier number;
        version (0 | 1 | automatic);
    }
```

To specify the threshold for the adaptation of the detection time, include the threshold statement:

```
detection-time {
    threshold milliseconds;
}
```

To specify the minimum transmit and receive interval for failure detection, include the minimum-interval statement:

```
minimum-interval milliseconds;
```

NOTE: Specifying an interval smaller than 300 ms can cause undesired BFD flapping.

To specify the minimum receive interval for failure detection, include the minimum-receive-interval statement:

```
minimum-receive-interval milliseconds;
```

To specify the detection time multiplier for failure detection, include the multiplier statement:

```
multiplier number;
```

To specify the threshold for detecting the adaptation of the transmit interval, include the threshold statement:

```
transmit-interval {
    threshold milliseconds;
}
```

The threshold value must be greater than the transmit interval.

To specify only the minimum transmit interval for failure detection, include the minimum-interval statement:

```
transmit-interval {
   minimum-interval milliseconds;
}
```

To specify the BFD version used for detection, include the version statement:

```
version (0 | 1 | automatic);
```

For a list of hierarchy levels at which you can configure these statements, see the statement summary sections for these statements.

Disabling Strict Address Check

If the sender of a RIP message does not belong to the subnet of the interface, the message is discarded. This situation may cause problems with dropped packets when RIP is running on point-to-point interfaces, or when the addresses on the interfaces do not fall in the same subnet. You can resolve this by disabling strict address checks on the RIP traffic.

To disable strict address checks, include the any-sender statement:

```
any-sender;
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

NOTE: The any-sender statement is supported only for peer-to-peer interfaces.

Tracing RIP Protocol Traffic

To trace RIP protocol traffic, you can specify options in the global traceoptions statement at the [edit routing-options] hierarchy level, and you can specify RIP-specific options by including the traceoptions statement:

```
traceoptions {
    file name <replace> <size size> <files number> <no-stamp>
      <(world-readable | no-world-readable)>;
    flag flag <flag-modifier> <disable>;
}
```

For a list of hierarchy levels at which you can configure this statement, see the statement summary section for this statement.

You can specify the following RIP-specific options in the RIP traceoptions statement:

- auth—Trace RIP authentication.
- error—Trace RIP errors.
- expiration—Trace RIP route expiration processing.
- holddown—Trace RIP hold-down processing.
- packets—Trace all RIP packets.

- request—Trace RIP information packets.
- trigger—Trace RIP triggered updates.
- update—Trace RIP update packets.

NOTE: Use the traceoption flags detail and all with caution. These flags may cause the CPU to become very busy.

For general information about tracing and global tracing options, see "Tracing Global Routing Protocol Operations" on page 118.

Example: Tracing RIP Protocol Traffic

Trace only unusual or abnormal operations to /var/log/routing-log, and trace detailed information about all RIP packets to /var/log/rip-log:

```
[edit]
routing-options {
    traceoptions {
        file /var/log/routing-log;
        flag errors;
    }
}
protocols {
    rip {
        traceoptions {
            file /var/log/rip-log;
            flag packets detail;
        }
    }
}
```

Configuring RIP

Configure RIP (for routing instances, include the statements at the [edit routing-options *routing-instance-name* protocols rip] hierarchy level):

```
[edit policy-options]
policy-statement redist-direct {
    from protocol direct;
    then accept;
}
```

```
[edit]
interfaces {
    so-0/0/0 {
        unit 0 {
            inet;
        }
    }
    at-1/1/0 {
        unit 0 {
            inet;
        }
    }
    at-1/1/0 {
        unit 42 {
            inet;
        }
    }
    at-1/1/1 {
        unit 42 {
            inet;
        }
    }
}
policy-statement redist-direct {
    from protocol direct;
    then accept;
}
[edit protocols rip]
metric-in 3;
receive both;
group wan {
    metric-out 2;
    export redist-direct;
    neighbor so-0/0/0.0;
    neighbor at-1/1/0.0;
    neighbor at-1/1/0.42;
    neighbor at-1/1/1.42 {
        receive version-2;
    }
}
group local {
    neighbor ge-2/3/0.0 {
        metric-in 1;
        send broadcast;
    }
}
```