

Chapter 11

RSVP Configuration Guidelines

To configure RSVP, you include statements at the [edit protocols rsvp] hierarchy level of the configuration.

```
protocols {
  rsvp {
    disable;
    keep-multiplier number;
    refresh-time seconds;
    traceoptions {
      file filename <replace> <size size> <files number> <no-stamp>
        <(world-readable | no-world-readable)>;
      flag flag <flag-modifier> <disable>;
    }
    interface interface-name {
      disable;
      (aggregate | no-aggregate);
      authentication-key key;
      bandwidth bps;
      hello-interval seconds;
      subscription percentage;
    }
  }
}
```

By default, RSVP is disabled.

This chapter describes the minimum required configuration and discusses the following tasks for configuring RSVP:

- Enable RSVP on page 118
- Configure RSVP Aggregation on page 119
- Configure the RSVP Hello Interval on page 119
- Configure RSVP Authentication on page 120
- Reserve Bandwidth on an Interface on page 120
- Configure RSVP Timers on page 121
- Trace RSVP Protocol Traffic on page 121
- Configure RSVP and MPLS on page 123

Minimum RSVP Configuration

To enable RSVP on all interfaces, include the following statement in the configuration file. All other RSVP configuration statements are optional.

```
[edit]
protocols {
  rsvp {
    interface all;
  }
}
```

Enable RSVP

To enable RSVP, including the following statements at the [edit] hierarchy level:

```
[edit]
protocols {
  rsvp {
    interface interface-name;
  }
}
```

To enable RSVP on all interfaces, specify *all* for *interface-name*.

If you have configured interface properties on a group of interfaces and want to disable RSVP on one of the interfaces, include the *disable* statement within the *rsvp* interface statement:

```
[edit]
protocols {
  rsvp {
    interface interface-name {
      disable;
    }
  }
}
```

Configure RSVP Aggregation

The resource requirements—processing, bandwidth, and memory—for running RSVP on a router increase proportionally with the number of sessions. Handling large numbers of refresh messages transmitted between RSVP neighbors is crucial for supporting large numbers of sessions. Path and Resv messages typically represent the majority of refreshes.

If topology failures occur, every node adjacent to the failure might want to notify all affected sender and receiver nodes. These notification messages are either tear or error messages, and they typically represent a flood that ripples out from the original failure point.

Aggregation provides a mechanism for reducing message flooding and network overload. It also enhances the efficiency and reliability in delivering RSVP tear or error messages. Note that RSVP aggregation is called "Bundle Message" in the internet draft RSVP Refresh Reduction Extensions.

By default, aggregation is disabled on all interfaces. For interoperability with other routers, you might need to keep aggregation disabled. However, we recommend that you enable aggregation between Juniper Networks routers to improve scalability. If you have several thousand MPLS LSPs, you must enable aggregation to ensure stable operation. To enable aggregation, include the aggregate statement at the [edit protocols rsvp interface *interface-name*] hierarchy level:

```
[edit protocols rsvp interface interface-name]  
aggregate;
```

Configure the RSVP Hello Interval

RSVP hello packets enable RSVP nodes to detect the loss of a neighboring node's RSVP state information. (Losses typically occur when the neighboring router restarts or the link fails.) In standard RSVP, such detection occurs as a consequence of RSVP's soft-state model. However, detection typically requires several minutes to time out the soft state. RSVP hello packets detect the neighboring node's state changes much more quickly, usually within 10-20 seconds.

Between hello-capable neighbors, hello packets are sent unicast toward each other. A loss of (2 x keep-multiplier + 1) consecutive hello packets causes the neighbor's state to go down, and all RSVP sessions to and from that neighbor are declared to be down.

JUNOS RSVP hello packets are optional and are backwards compatible with RSVP implementations that do not support hello packets. For neighbors that do not support hello packets, RSVP uses the soft-state timeout for loss detection.

If all neighboring nodes support hello packets, you can reduce the refresh overhead (by increasing the value set in the refresh-time statement) without adversely affecting the node or link failure detection time. Also, the network can scale to a larger number of sessions because the refresh operations consume less CPU and bandwidth. For information about setting the refresh overhead, see "Configure RSVP Timers" on page 121.

By default, RSVP sends hello packets every 3 seconds. To modify how often RSVP sends hello packets, include the hello-interval statement at the [edit protocols rsvp interface *interface-name*] hierarchy level:

```
[edit protocols rsvp interface interface-name]  
hello-interval seconds;
```

Configure RSVP Authentication

All RSVP protocol exchanges can be authenticated to guarantee that only trusted neighbors participate in setting up reservations. By default, RSVP authentication is disabled.

RSVP authentication uses an HMAC-MD5 message-based digest. This scheme produces a message digest based on a secret authentication key and the message contents. (The message contents also include a sequence number.) The computed digest is transmitted with RSVP messages. Once you have configured authentication, all received and transmitted RSVP messages with all neighbors are authenticated on this interface.

MD5 authentication also provides protection against forgery and message modification. However, it does not provide confidentiality because all messages are sent in clear text, and it does not prevent replay attacks.

By default, authentication is disabled. To enable authentication, configure a key on each interface by including the authentication-key statement at the [edit protocols rsvp interface *interface-name*] hierarchy level:

```
[edit protocols rsvp interface interface-name]
  authentication-key key;
```

Reserve Bandwidth on an Interface

For each interface on which RSVP is enabled, by default, RSVP permits all the interface's bandwidth (100 percent) to be used for RSVP reservations.

Oversubscription on an interface occurs when the aggregate demand of all RSVP sessions is allowed to exceed physical capacity of the link. You can use oversubscription to take advantage of the statistical nature of traffic patterns and to permit higher utilization of links. In particular, you can use oversubscription in places where peak utilizations of traffic do not coincide in time.

Undersubscription on an interface occurs when the total demand of all RSVP sessions is always less than the physical capacity of the link. You can use undersubscription to bound utilization of links and reduce congestion.

You can modify the link bandwidth used for RSVP reservations, either decreasing it below 100 percent or oversubscribing the interface. To do this, include the subscription statement at the [edit protocols rsvp interface *interface-name*] hierarchy level:

```
[edit protocols rsvp interface interface-name]
  subscription percentage;
```

percentage is the percentage of the interface's bandwidth that RSVP allows to be used for reservations. It can be a value from 0 through 65000 percent. If you specify a value greater than 100, you are oversubscribing the interface.

You can use the subscription factor to shut down new RSVP sessions on a per-interface basis. If you set the percentage to 0, no new sessions (including those with zero bandwidth requirements) are permitted on the interface. Existing RSVP sessions are not affected by changing the subscription factor. To clear an existing session, issue the clear rsvp session command.

Configure RSVP Timers

RSVP uses two interrelated timing parameters:

The refresh time controls the interval between the successive generation of refresh messages. Refresh messages include Path and Resv messages. Refresh messages are sent periodically so that reservation states in neighboring nodes do not time out. Each node chooses a value for the refresh timer independently. Each Path and Resv message carries the refresh timer value, and the receiving node extracts this value from the messages.

The keep multiplier is a locally configured small integer in the range 1 through 255.

To determine the lifetime of a reservation state, use the following formula:

$$lifetime = (keep-multiplier + 0.5) * 1.5 * refresh-time$$

In the worst case, $(keep-multiplier - 1)$ successive refresh messages must be lost before a reservation state is deleted.

By default, the refresh timer value is 30 seconds. To modify this value, include the refresh-time statement at the [edit protocols rsvp] hierarchy level:

```
[edit protocols rsvp]
refresh-time seconds;
```

The default value of the keep multiplier is 3. To modify this value, include the keep-multiplier statement at the [edit protocols rsvp] hierarchy level:

```
[edit protocols rsvp]
keep-multiplier number;
```

Trace RSVP Protocol Traffic

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

```
[edit protocols rsvp]
traceoptions {
  file filename <replace> <size size> <files number> <no-stamp>
  <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

You can specify the following RSVP-specific flags in the RSVP traceoptions statement:

error—Trace all detected error conditions.

packets—Trace all RSVP messages, including Path, Resv, PathTear, ResvTear, PathErr, ResvErr, and ResvConf messages.

path—Trace Path messages.

pathtear—Trace PathTear messages.

resv—Trace Resv messages.

resvtear—Trace ResvTear messages.

state—Trace session state transitions.

For general information about tracing and global tracing options, see the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

Examples: Trace RSVP Protocol Traffic

Trace RSVP Path messages in detail:

```
[edit]
protocols {
  rsvp {
    traceoptions {
      file rsvp size 10m files 5;
      flag path;
    }
  }
}
```

Trace all RSVP messages:

```
[edit]
protocols {
  rsvp {
    traceoptions {
      file rsvp size 10m files 5;
      flag packets;
    }
  }
}
```

Trace all RSVP error conditions:

```
[edit]
protocols {
  rsvp {
    traceoptions {
      file rsvp size 10m files 5;
      flag error;
    }
  }
}
```

Configure RSVP and MPLS

The primary purpose of the JUNOS RSVP software is to support dynamic signaling within LSPs. When you enable both MPLS and RSVP on a router, MPLS becomes a client of RSVP. No additional configuration is required to bind MPLS and RSVP.

You can configure MPLS to set up signaled paths using the `label-switched-path` statement. Each LSP translates into a request for RSVP to initiate an RSVP session. This request is passed through the internal interface between label switching and RSVP. After examining the request information, checking RSVP states and the local routing tables, RSVP initiates one session for each LSP. The session is sourced from the local router and is destined to the target of the LSP.

When an RSVP session is successfully created, the LSP is set up along the paths created by the RSVP session. If the RSVP session is unsuccessful, RSVP notifies MPLS of its status. It is up to MPLS to try to initiate backup paths or to continue retrying the initial path.

To pass label-switching signaling information, RSVP supports four additional objects: Label Request Object, Label Object, Explicit Route Object, and Record Route Object. For an LSP to be set up successfully, all routers along the path must support MPLS, RSVP, and these four objects. Of the four objects, Record Route Object is not mandatory.

To configure MPLS and make it a client of RSVP, do the following:

Enable MPLS on all routers that are to participate in label switching (that is, on all routers that might be part of a label-switching path).

Enable RSVP on all routers and on all router interfaces that form the LSP.

Configure the routers that are to be the beginning of the LSP.

Example: Configure RSVP and MPLS

The following shows an example configuration for a router at the beginning of an LSP:

```
[edit]
protocols {
  mpls {
    label-switched-path sf-to-london {
      to 192.168.1.4;
    }
  }
  rsvp {
    interface so0-0/0/0;
  }
}
```

The following shows an example configuration for all the other routers that form the LSP:

```
[edit]
protocols {
  mpls {
    interface so0-0/0/0;
  }
  rsvp {
    interface so0-0/0/0;
  }
}
```

}