

Chapter 5

Configure MPLS Signaled LSPs

To configure MPLS signaled LSPs, you create an LSP that runs from the ingress router to the egress router. (For information on LDP signaled LSPs, see “Configure LDP” on page 143.) To create the LSP, you configure only the ingress router; you do not have to configure any other routers. You can configure the LSP so that the JUNOS software makes all the forwarding decisions, or you can configure some or all of the routers in the path. The LSP is set up by RSVP, through RSVP signaling messages. The JUNOS software automatically negotiates, assigns, releases, and reuses labels. Automatically assigned labels have a value from 1024 through 1048575.

To configure signaled LSPs across a network, perform the following tasks:

- Configure the Ingress Router for Signaled LSPs on page 39

- Configure All Other MPLS Routers for Signaled LSPs on page 64

- Enable RSVP on page 64

For a configuration example, see “Examples: Configure Signaled LSPs” on page 64.

Configure the Ingress Router for Signaled LSPs

To configure signaled LSPs, perform the following tasks on the ingress router:

- Create a Named Path on page 39

- Create an LSP on page 41

Create a Named Path

To configure signaled LSPs, you must first create one or more named paths on the ingress router. For each path, you can specify some or all of the transit routers in the path, or you can leave it empty.

Each path name can be up to 16 characters and can contain letters, digits, periods, and hyphens. The name must be unique within the ingress router. Once a named path is created, you can configure LSPs using the named path on the primary or on the secondary statement at the [edit protocols mpls label-switched-path *label-path-name*] hierarchy level. You can specify the same named path on any number of LSPs.

To create an empty path, create a named path by including the following form of the path statement at the [edit protocols mpls] hierarchy level. This form of the path statement is empty, which means that any path between the ingress and egress routers is accepted. In actuality, the path used tends to be the same path as is followed by destination-based, best-effort traffic.

```
[edit protocols mpls]
path path-name;
```

To create a path in which you specify some or all of the transit routers in the path, include the following form of the path statement at the [edit protocols mpls] hierarchy level, specifying one *address* for each transit router:

```
[edit protocols mpls]
path path-name {
  address | host name <strict | loose>;
}
```

In this form of the path statement, you specify one or more transit router addresses. Specifying the ingress and/or egress routers is optional. You can specify the address or host name of each transit router, although you do not need to list each transit router if its type is loose. Specify the addresses in order, starting with the ingress router (optional) or the first transit router, and continuing sequentially along the path up to the egress router (optional) or the router immediately before the egress router. You need to specify only one address per router hop. If you specify more than one address for the same router, only the first address is used; the additional addresses are quietly ignored and truncated.

For each router address, you specify the type, which can be one of the following:

strict—(Default) The route taken from the previous router to this router is a direct path and cannot include any other routers. If *address* is an interface address, this router also ensures that the incoming interface is the one specified. Doing this is useful when there are parallel links between the previous router and this router. It also ensures that routing can be enforced on a per-link basis.

For strict addresses, you must ensure that the router immediately preceding the router you are configuring has a direct connection to that router. The address can be a loopback interface address, in which case the incoming interface is not checked.

loose—The route taken from the previous router to this router need not be a direct path and can include other routers and can be received on any interface. The address can be any interface address or the address of the loopback interface.

Examples: Create a Named Path

The following path, `to-hastings`, specifies the complete strict path from the ingress to the egress routers through 14.1.1.1, 13.1.1.1, 12.1.1.1 and 11.1.1.1, in that order. There cannot be any intermediate routers except the ones specified. However, there can be intermediate routers between 11.1.1.1 and the egress router because the egress router is not specifically listed in the path statement. To prevent intermediate routers before egress, configure the egress router as the last router, with a strict type.

```
[edit protocols mpls]
path to-hastings {
  14.1.1.1 strict;
  13.1.1.1 strict;
  12.1.1.1 strict;
  11.1.1.1 strict;
}
```

The following path, `alt-hastings`, allows any number of intermediate routers between routers 14.1.1.1 and 11.1.1.1. In addition, intermediate routers are permitted between 11.1.1.1 and the egress router.

```
[edit protocols mpls]
path alt-hastings {
  14.1.1.1 strict;
  11.1.1.1 loose;
}
```

Create an LSP

The second step in configuring signaled LSPs is to create one or more LSPs and define the properties associated with the label-switched path on the ingress router. To configure an LSP, include the label-switched-path statement at the [edit protocols mpls] hierarchy level:

```
[edit protocols mpls]
label-switched-path lsp-path-name {
  disable;
  to address;
  from address;
  adaptive;
  admin-group {
    exclude group-names;
    include group-names;
  }
  bandwidth bps;
  class-of-service cos-value;
  fast-reroute {
    bandwidth bps;
    exclude group-names;
    hop-limit number;
    include group-names;
  }
  hop-limit number;
  ldp-tunneling;
  metric number;
  no-cspf;
  no-decrement-ttl;
  optimize-timer seconds;
  preference preference;
  priority setup-priority hold-priority;
```

```

    (random | least-fill | most-fill);
    (record | no-record);
    retry-limit number;
    retry-timer seconds;
    standby;
    primary path-name {
        adaptive;
        admin-group {
            exclude group-names;
            include group-names;
        }
        bandwidth bps;
        class-of-service cos-value;
        hop-limit number;
        no-cspf;
        optimize-timer seconds;
        preference preference;
        priority setup-priority hold-priority;
        (record | no-record);
        standby;
    }
    secondary path-name {
        adaptive;
        admin-group {
            exclude group-names;
            include group-names;
        }
        bandwidth bps;
        class-of-service cos-value;
        hop-limit number;
        no-cspf;
        optimize-timer seconds;
        preference preference;
        priority setup-priority hold-priority;
        (record | no-record);
        standby;
    }
}
]

```

Each LSP must have a name, *lsp-path-name*, which can be up to 32 characters long and can contain letters, digits, periods (.), and hyphens (-). The name must be unique within the ingress router. For ease of management and identification, configure unique names across the entire domain.

When you configure LSPs, you can specify the following statements either for each LSP or for each path. (You configure LSPs at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level, and you configure paths at the [edit protocols mpls label-switched-path *lsp-path-name* primary] or [edit protocols mpls label-switched-path *lsp-path-name* secondary] hierarchy level.) For statements that you configure on a per-LSP basis, the value applies to all paths in the LSP. For statements that you configure on a per-path basis, the path value overrides the per-LSP value.

```

adaptive
admin-group
bandwidth
class-of-service
hop-limit

```

no-cspf

optimize-timer

preference

priority

record or no-record

standby

For each LSP, you can configure the following properties:

Configure the Address of the Egress Router on page 44

Configure the Address of the Ingress Router on page 44

Configure the Primary and Secondary LSPs on page 45

Configure Fast Reroute on page 45

Configure Addresses to Associate with the LSP on page 48

Configure Path Connection Retry Information on page 49

Configure the Dynamic LSP Metric on page 49

Configure the Static LSP Metric on page 50

Configure CSPF Tie Breaking on page 50

Disable Normal TTL Decrementing on page 51

For each LSP and for each primary and secondary path, you can configure the following properties:

Disable Constrained Path LSP Computation on page 52

Configure Administrative Groups on page 53

Configure the LSP Preference on page 55

Configure Whether to Record Path Routes on page 55

Configure the MPLS CoS Value on page 55

Configure an LSP to be Adaptive on page 57

Configure Priority and Preemption on page 58

Optimize Signaled LSPs on page 59

Configure the Maximum Path Length on page 60

Configure the Path Bandwidth on page 60

- Configure the Standby State on page 60

- Configure LDP Tunneling on page 61

Configure the Address of the Egress Router

When configuring an LSP, you must specify the address of the egress router by including the `to` statement at the `[edit protocols mpls label-switched-path lsp-path-name]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-path-name]  
  to address;
```

When you are setting up an LSP, the `to` statement is the only required statement. All other statements are optional.

After the LSP is established, the address of the egress router is installed as a host route in the routing table. Then, this route can be used by BGP to forward traffic.

To have the software send BGP traffic over an LSP, the address of the egress router is the same as the address of the BGP next hop. You can specify the egress router's address as any one of the router's interface addresses or as the BGP router ID. If you specify a different address, even if the address is on the same router, BGP traffic is not sent over the LSP.

To determine the address of the BGP next hop, use the `show route detail` command. To determine the destination address of an LSP, use the `show mpls lsp` command. To determine whether a route has gone through an LSP, use the `show route` or `show route forwarding-table` command. In the output of these last two commands, the `label-switched-path` or `push` keyword included with the route indicates it has passed through an LSP. Also, use the `traceroute` command to trace the actual path that the route leads to. This is another indication as to whether a route has passed through an LSP.

You also can manipulate the address of the BGP next hop by defining a BGP import policy filter that sets the route's next-hop address.

Configure the Address of the Ingress Router

The local router always is considered to be the ingress router, which is the beginning of the LSP. The software automatically determines the proper outgoing interface and IP address to use to reach the next router in an LSP.

By default, the router ID is chosen as the address of the ingress router. To override the automatic selection of the source address, specify a source address in the `from` statement at the `[edit protocols mpls label-switched-path lsp-path-name]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-path-name]  
  from address;
```

The outgoing interface used by the LSP is not affected by the source address that you configure.

Configure the Primary and Secondary LSPs

By default, an LSP routes itself hop by hop toward the egress router. The LSP tends to follow the shortest path as dictated by the local routing table, usually taking the same path as destination-based, best-effort traffic. These paths are “soft” in nature because they automatically reroute themselves whenever a change occurs in a routing table or in the status of a node or link.

To configure the path so that it follows a particular route, create a named path using the path statement, as described in “Create a Named Path” on page 39. Then you apply the named path by including the primary or secondary statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level:

```
[edit protocols mpls label-switched-path lsp-path-name]
primary path-name {
  ...
}
secondary path-name {
  ...
}
```

A named path can be referenced by any number of LSPs.

The primary statement creates the primary path, which is the LSP’s preferred path. The secondary statement creates an alternative path. If the primary path can no longer reach the egress router, the alternative path is used.

When the software switches from the primary to a secondary path, it continuously attempts to revert to the primary path, switching back to it when it is again reachable, but no sooner than the retry time specified in the retry-timer statement. (For more information, see “Configure Path Connection Retry Information” on page 49.)

You can configure zero or one primary path. If you do not configure a primary path, the first secondary path that is established is selected as the path.

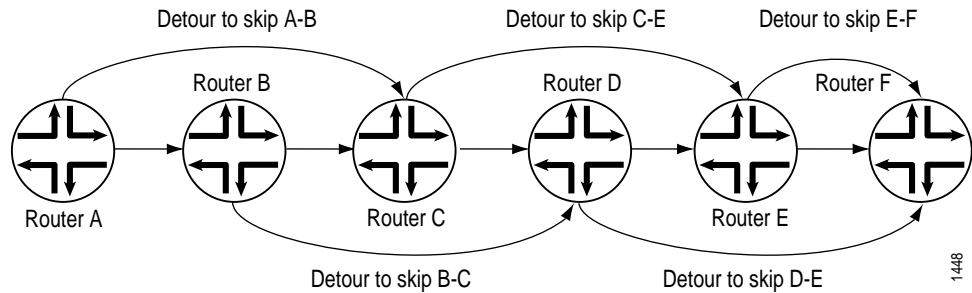
You can configure zero or more secondary paths. All secondary paths are equal, and the software tries them in the order that they are listed in the configuration. The software does not attempt to switch among secondary paths. If the current secondary path is not available, the next one is tried. To create a set of equal paths, specify secondary paths without specifying a primary path.

If you do not specify any named paths, or if the path that you specify is empty, the software makes all routing decisions necessary to reach the egress router.

Configure Fast Reroute

Fast reroute provides a mechanism for automatically rerouting traffic on an LSP if a node or link in an LSP fails, thus reducing the loss of packets traveling over the LSP. Fast rerouting is accomplished by precomputing and pre-establishing a number of detours along the LSP. Figure 13 illustrates an LSP from Router A to Router F, showing some of the detours that are established for the LSP. Each detour is established by an upstream node with the intent of avoiding the link toward the immediate downstream node and the immediate downstream node itself. Each detour might traverse through one or more label-switched routers that are not shown in the figure.


Figure 13: Detours Established for an LSP using Fast Reroute



If a node detects either that a downstream link has failed (using a link-layer specific liveness detection mechanism) or that a downstream node has failed (for example, using the RSVP neighbor hello protocol), the node quickly splices the traffic onto the detour and, at the same time, signals the ingress router about the link or node failure. Figure 14 illustrates the detour taken when the link between Router B and Router C fails.

If the network topology is not rich enough, some of the detour might not succeed. For example, the detour from Router A to Router C in Figure 13 cannot traverse link A-B and Router B. If such a path is not possible, the detour does not occur.

Figure 14: Detour after the Link from Router B to Router C Fails



The time required for a fast-rerouting detour to take effect depends on two independent time intervals:

Amount of time to detect that there is a link or node failure—This time interval depends greatly on the link layer in use and the nature of the failure. For example, failure detection on an SDH/SONET link typically is much faster than on a Gigabit Ethernet link, and both are much faster than detection of a router failure.


Amount of time required to splice the traffic onto the detour—This time interval is primarily the CPU time required to update the routing table and then to update the forwarding table. The amount of time depends on the current CPU load and how busy the other routing protocols are that are sharing the CPU.

Fast reroute is a short-term patch to reduce packet loss. Because detour computation might not reserve adequate bandwidth, the detours might introduce congestion on the alternate links. The ingress router is the only router that is fully aware of LSP policy constraints and, therefore, is the only router able to come up with adequate long-term alternate paths.

Detours are created using RSVP and, like all RSVP sessions, they require extra state and overhead in the network. For this reason, each node establishes at most one detour for each LSP that has fast reroute enabled. Creating more than one detour for each LSP increases the overhead, but serves no practical purpose.

To reduce network overhead further, each detour attempts to merge back into the LSP as soon as possible after the failed node or link. If you can consider an LSP that travels through N router nodes, it is possible to create $N - 1$ detours. For instance, in Figure 15, the detour tries to merge back into the LSP at Router D instead of at Router E or Router F. Merging back into the LSP makes the detour scalability problem more manageable. If topology limitations prevent the detour from quickly merging back into the LSP, detours merge with other detours automatically.

Figure 15: Detours Merging into Other Detours



Fast reroute protects traffic against any single point of failure between the ingress and egress routers. If there are multiple failures along an LSP, it is possible that fast reroute itself might fail. Also, fast reroute does not protect against failure of the ingress or egress routers.

Computing and setting up detours are done independently at each node. On a node, if an LSP has fast reroute enabled and if a downstream link or node can be identified, the router performs a CSPF computation using the information in the local traffic engineering database (TED). For this reason, detours rely on your IGP's supporting traffic engineering extensions. Without the TED, detours cannot be established.

Detour computations might not succeed the first time. If a computation fails, the router recomputes detours approximately once every refresh interval until the computation succeeds.

To enable fast reroute on an LSP, include the `fast-reroute` statement at the `[edit protocols mpls label-switched-path lsp-path-name]` hierarchy level on the ingress router:

```
[edit protocols mpls label-switched-path lsp-path-name]
fast-reroute {
  bandwidth bps;
  exclude group-names;
  hop-limit number;
  include group-names;
}
```

You do not need to configure fast reroute on the LSP's transit and egress routers. Once fast reroute is enabled, the ingress router signals all the downstream routers that fast reroute is enabled on the LSP, and each downstream router does its best to set up detours for the LSP. If a downstream router does not support fast reroute, it ignores the request to set up detours and continues to support the LSP. A router that does not support fast reroute will cause some of the detours to fail, but otherwise has no impact on the LSP.

By default, no bandwidth is reserved for the rerouted path. To allocated bandwidth for the rerouted path, include the `bandwidth` statement. The bandwidth does not need to be identical to that allocated for the LSP.

Hop-limit constraints define how many more routers a detour is allowed to traverse compared to the LSP itself. By default, the hop limit is set to 6. For example, if an LSP traverses four routers, any detour for the LSP can be up to 10 (that is, 4 + 6) router hops, including the ingress and egress routers.

By default, a detour inherits the same administrative (coloring) group constraints as its parent LSP when CSPF is determining the alternate path. Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. If you specify the include statement when configuring the parent LSP, all links traversed by the alternate session must have at least one color found in the list of groups. If you specify the exclude statement when configuring the parent LSP, all links must not have a color found in the list of groups. For more information about administrative group constraints, see “Configure Administrative Groups” on page 53.

Configure Addresses to Associate with the LSP

By default, a host route toward the egress router is installed in the inet.3 routing table. (The host route address is the one you configure in the to statement.) Installing the host route allows BGP to perform next-hop resolution. It also prevents the host route from interfering with prefixes learned from dynamic routing protocols and stored in the inet.0 routing table.

Unlike the routes in the inet.0 table, routes in the inet.3 table are not copied to the Packet Forwarding Engine, and hence they cause no changes in the system forwarding table directly. You cannot ping or traceroute through these routes. The only use for inet.3 is to permit BGP to perform next-hop resolution. To examine the inet.3 table, use the show route table inet.3 command.

To inject additional routes into the inet.3 routing table, include the install statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level:

```
[edit protocols mpls label-switched-path lsp-path-name]
install {
    destination/mask <active>;
}
```

The specified routes are installed as aliases into the routing table when the LSP is established. Installing additional routes allows BGP to resolve next hops within the specified prefix and to direct additional traffic for these next hops to a particular LSP.

Including the active option with the install statement installs the specified prefix into the inet.0 routing table, which is the primary forwarding table. The result is a route that is installed in the forwarding table any time the LSP is established, which means you can ping or traceroute the route. Use this option with care, because this type of prefix is very similar to a static route.

You use alias routes for routers that have multiple addresses being used as BGP next hops, or for routers that are not MPLS-capable. In either of these cases, the LSP can be configured to another MPLS-capable system within the local domain, which then acts as a “border” router. The LSP then terminates on the border router and, from that router, Layer 3 forwarding takes the packet to the true next-hop router.

In the case of an interconnect, the domain’s border router can act as the proxy router and can advertise the prefix for the interconnect if the border router is not setting the BGP next hop to itself.

In the case of a POP that has routers that do not support MPLS, one router (for example, a core router) that supports MPLS can act as a proxy for the entire POP and can inject a set of prefixes that cover the POP. Thus, all routers within the POP can advertise themselves as IBGP next hops, and traffic can follow the LSP to reach the core router. This means that normal IGP routing would prevail within the POP.

You cannot use the ping or traceroute commands on routes in the inet.3 routing table.

For BGP next-hop resolution, it makes no difference whether a route is in inet.0 or inet.3. The route with the best match (longest mask) is chosen. Among multiple best-match routes, the one with the highest preference value is chosen.

Configure Path Connection Retry Information

The ingress router might make many attempts to connect and reconnect to the egress router using the primary path. You can control how often the ingress router tries to establish a connection using the primary path and how long it waits between retry attempts.

The retry timer configures how long the ingress router waits before trying to connect again to the egress router using the primary path. The default retry time is 30 seconds. The time can be from 1 through 600 seconds. To modify this value, include the `retry-timer` statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level:

```
[edit protocols mpls label-switched-path lsp-path-name]  
  retry-timer seconds;
```

By default, no limit is set to the number of times an ingress router attempts to establish or re-establish a connection to the egress router using the primary path. To limit the number of attempts, include the `retry-limit` statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level:

```
[edit protocols mpls label-switched-path lsp-path-name]  
  retry-limit number;
```

The limit can be a value up to 10000. When the retry limit is exceeded, no more attempts are made to establish a path connection. At this point, intervention is required to restart the primary path.

If you set a retry limit, it is reset to 1 each time a successful primary path is created.

Configure the Dynamic LSP Metric

If no specific metric is configured, an LSP attempts to track the IGP metric toward the same destination (the to address of the LSP). IGP includes OSPF, IS-IS, RIP, and static routes. BGP and other RSVP/LDP routes are excluded.

For example, if the OSPF metric toward a router is 20, then all LSPs toward that router will automatically inherit metric 20. If the OSPF toward a router later changes to a different value, all LSP metrics change accordingly. If there are no IGP routes toward the router, the LSP will raise its metric to 65535.

Note that in this case, the LSP metric is completely determined by IGP; it bears no relationship to the actual path the LSP is currently traversing through. If LSP reroutes (such as through reoptimization), its metric doesn't change, and thus it remains transparent to users. Dynamic metric is the default behavior; no configuration is required.

Configure the Static LSP Metric

You can manually assign a fixed metric value to an LSP. Once configured using the metric statement at the [edit protocols mpls label-switched-path lsp-name] hierarchy level, the LSP metric is fixed and will not change:

```
[edit protocols mpls label-switched-path lsp-name]
metric number;
```

The LSP metric has several uses:

When there are parallel LSPs with the same egress router, the metrics are compared to see which LSP has the lowest metric value (the lowest cost) and therefore the preferred path to the destination. If the metrics are the same, the traffic is shared.

Adjusting the metric values can force traffic to prefer some LSPs over others, regardless of the underlying IGP metric.

When an IGP shortcut is enabled (see “IGP Shortcuts” on page 24), an IGP route might be installed in the routing table with an LSP as the next hop, if the LSP is on the shortest path to the destination. In this case, the LSP metric is added to the other IGP metrics to determine the total path metric. For example, if an LSP whose ingress router is X and egress router is Y is on the shortest path to destination Z, the LSP metric is added to the metric for the IGP route from Y to Z to determine the total cost of the path. If several LSPs are potential next hops, the total metrics of the paths are compared to determine which path is preferred (that is, has the lowest total metric). Or, IGP paths and LSPs leading to the same destination could be compared using the metric value to determine which path is preferred.

By adjusting the LSP metric, you can force traffic to prefer LSPs, to prefer the IGP path, or to share the load among them.

If router X and Y are BGP peers, and if there is an LSP between them, the LSP metric represents the total cost to reach Y from X. If for any reason the LSP reroutes, the underlying path cost might change significantly, but X’s cost to reach Y remains the same (the LSP metric), which allows X to report through BGP MED a stable metric to downstream neighbors. As long as Y remains reachable through the LSP, no changes are visible to downstream BGP neighbors.

Configure CSPF Tie Breaking

When CSPF is selecting a path for an LSP, if there are several equal-cost paths there is a tie-breaking process used. For information about how CSPF selects a path, see “How CSPF Selects a Path” on page 23. To configure a random tie-breaking rule for CSPF to use to choose among equal-cost paths, include the random statement at the [edit protocols mpls path label-switched-path lsp-path-name]:

```
[edit protocols mpls path label-switched-path lsp-path-name]
random;
```

To prefer the path with the least utilized links, include the least-fill statement at the [edit protocols mpls path label-switched-path lsp-path-name]:

```
[edit protocols mpls path label-switched-path lsp-path-name]
least-fill;
```

To prefer the path with the most utilized links, include the most-fill statement at the [edit protocols mpls path label-switched-path *lsp-path-name*]:

```
[edit protocols mpls path label-switched-path lsp-path-name]
most-fill;
```

Configure Load Balancing LSPs Without CSPF

LSP tends to load-balance its placement by randomly selecting one of the equal-cost next hops and using it exclusively. The random selection is made independently at each transit router, and is made by comparing IGP metrics alone. No consideration is given to bandwidth or congestion levels.

Disable Normal TTL Decrementing

By default, the TTL field value in the packet header is decremented by 1 for every hop the packet traverses in the LSP, thereby preventing loops. If the TTL field value reaches 0, packets are dropped and an ICMP error packet might be sent to the originating router.

If normal TTL decrement is disabled, the TTL field of IP packets entering LSPs are decremented by only 1 upon transiting the LSP, making the LSP appear as a one-hop router to diagnostic tools, such as traceroute. This is done by the ingress router, which pushes a label on IP packets with the TTL field in the label initialized to 255. The label's TTL field value is decremented by 1 for every hop the MPLS packet traverses in the LSP. On the penultimate hop of the LSP, the router pops the label, but does not write the label's TTL field value to the IP packet's TTL field. Instead, when the IP packet reaches the egress router, the IP packet's TTL field value is decremented by 1.

When you use traceroute to diagnose problems with an LSP, traceroute sees the ingress router, although the egress router performs the TTL decrement. Note that this assumes that traceroute is initiated outside of the LSP. The behavior of the traceroute is different if it is initiated from the ingress router of the LSP. In this case, the egress router would be the first router to respond to traceroute.

You can disable normal TTL decrementing in an LSP so that the TTL field value would not reach 0 before the packet reaches its destination, thus preventing the packet from being dropped. You can also disable normal TTL decrementing to make the MPLS cloud appear as a single hop, thereby hiding the network topology.

There are two ways to disable TTL decrementing:

On the ingress of the LSP, if you include the no-decrement-ttl statement at the [edit protocols mpls label-switched-path *lsp-path-name*] hierarchy level, the ingress router negotiates with all downstream routers using a proprietary RSVP object, to ensure all routers are in agreement. If negotiation succeeds, the whole LSP behaves as 1 hop to transit IP traffic.

```
[edit protocols mpls label-switched-path lsp-path-name]
no-decrement-ttl;
```

Note that the RSVP object is proprietary to JUNOS software, and might not work with other vendor software. Further, this potential incompatibility only applies to RSVP signaled LSPs, not LDP signaled LSPs. When you include the no-decrement-ttl statement, TTL hiding can be enforced on a per-LSP basis.

On the router, you can include the `no-propagate-ttl` statement at the `[edit protocols mpls]` hierarchy level. This configuration statements to all LSPs, regardless of whether they are RSVP-sigaled or LDP-sigaled. Once set, all future LSPs traversing through this router behave as a single hop to IP packets. LSPs established you enable this feature are not affected.

```
[edit protocols mpls]
no propagate-ttl;
```

If you include the `no-propagate-ttl` statement, make sure all routers are configured consistently within an MPLS domain; failing to do so might cause IP TTL to increase while in transit within LSPs. This can happen, for example, when the ingress router has `no-propagate-ttl` enabled but not the penultimate router, so the penultimate writes MPLS TTL (which starts from the ingress as 255) into the IP.

The operation of the `no-propagate-ttl` statement is more interoperable with other vendors' equipment. However, you must ensure all routers are configured identically.

Disable Constrained Path LSP Computation

If the IGP is a link state protocol and if it supports extensions that allow the current bandwidth reservation on each router's link to be reported, constrained path LSPs are computed by default.

The JUNOS implementations of IS-IS and OSPF include the extensions that support constrained-path LSP computation. In IS-IS, these extensions are enabled by default. (To disable this support, include the `disable` statement at the `[edit protocols isis traffic-engineering]` hierarchy level, as discussed in the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*. In OSPF, these extensions are disabled by default. To enable this support, include the `traffic-engineering` statement in the configurations of all routers running OSPF, as described in the *JUNOS Internet Software Configuration Guide: Routing and Routing Protocols*.

If IS-IS is enabled on a router or if you enable OSPF traffic engineering extensions, MPLS performs the constrained-path LSP computation by default.

Constrained-path LSP computation works as follows: LSPs advertise their link information in the IGP's link-state packets. These packets are flooded throughout the network and hence provide information to all nodes. This link information is placed into the traffic engineering database (TED) and provides each ingress router with LSP topology information and recent LSP bandwidth reservation information. When computing complete paths for LSPs, the ingress router uses the information in the TED, along with the requirements you configure for the LSP, including bandwidth (configured with the `bandwidth` statement), hop limit (configured with the `hop-limit` statement), and the address of the egress router (configured with the `to` statement).

Constrained-path LSPs have a greater chance of being established quickly and successfully for several reasons:

- The LSP computation takes into account the current bandwidth reservation.

- Constrained-path LSPs reroute themselves away from node failures and congestion.

When constrained-path LSP computation is enabled, you can configure the LSP so that it is periodically re-optimized, as described in "Optimize Signaled LSPs" on page 59.

When an LSP is being established or when an existing LSP fails, the constrained-path LSP computation is repeated periodically at the interval specified by the retry timer, until the LSP is set up successfully. Once the LSP is set up, no recomputation is done. For more information about the retry timer, see “Configure Path Connection Retry Information” on page 49.

By default, constrained-path LSP computation is enabled. You might want to disable constrained-path LSP computation when all nodes do not support the necessary traffic engineering extensions. To disable constrained-path LSP computation, include the `no-cspf` statement at the `[edit protocols mpls label-switched-path lsp-path-name]` or `[edit protocols mpls label-switched-path lsp-path-name (primary | secondary)]` hierarchy level:

```
no-cspf;
```

Configure Administrative Groups

Administrative groups, also known as link coloring or resource class, are manually assigned attributes that describe the “color” of links, such that links with the same color conceptually belong to the same class. You can use administrative groups to implement a variety of policy-based LSP setups.

Administrative groups are meaningful only when constrained-path LSP computation is enabled.

Administrative groups require three levels of configuration. First, configure a table of group names at the `[edit protocols mpls]` hierarchy level:

```
[edit protocols mpls]
admin-groups{
  group-name group-value;
}
```

You can assign up to 32 names and values (in the range 0 through 31), which define a series of names and their corresponding values. The administrative names and values must be identical across all routers within a single domain.

To configure administrative groups, follow these steps:

1. Define multiple levels of service quality:

```
[edit]
protocols {
  mpls {
    admin-groups {
      best-effort 1;
      copper 2;
      silver 3;
      gold 4;
      violet 5;
    }
  }
}
```

2. Define administrative groups for an interface. These groups serve to identify the administrative groups to which an interface belongs. You can assign multiple groups to an interface.

```
[edit]
protocols {
  mpls {
    interface interface name {
      admin-group [ group-name group-name...];
    }
  }
}
```

If you do not include the `admin-group` statement, an interface does not belong to any group.

IGPs use the group information to build link-state packets, which are then flooded throughout the network, providing information to all nodes in the network. At any router, the IGP topology, as well as administrative groups of all the links, is available.

Changing the interface's administrative group affects only new LSPs. Existing LSPs on the interface are not preempted or recomputed to keep the network stable. If LSPs need to be removed because of a group change, issue the `clear rsvp session` command.

3. Configure an administrative group constraint for each LSP or for each primary or secondary LSP path, at the `[edit protocols mpls label-switched-path lsp-path-name]` or `[edit protocols mpls label-switched-path lsp-path-name (primary | secondary)]` hierarchy level:

```
[edit]
protocols {
  mpls {
    label-switched-path lsp-path-name {
      to address;
      ...
      primary path-name {
        admin-group {
          exclude [ group-name group-name ... ];
          include [ group-name group-name ... ];
        }
      }
      secondary path-name {
        admin-group {
          exclude [ group-name group-name ... ];
          include [ group-name group-name ... ];
        }
      }
      admin-group {
        exclude [ group-name group-name ... ];
        include [ group-name group-name ... ];
      }
    }
  }
}
```

If you omit the include or exclude statements, the path computation proceeds unchanged using constrained-path LSP computation. If you configure an exclude list, all the links that are chosen must not have a color listed in the exclude list. If you configure an include list, all the links that are chosen must have at least one color found in the include list. Links that have no color are automatically disqualified by any include or exclude list.



Note

Changing the LSP's administrative group causes an immediate recomputation of the route; therefore, the LSP might be rerouted.

Configure the LSP Preference

As an option, you can configure multiple LSPs between the same pair of ingress and egress routers. This is useful for balancing the load among the LSPs because all LSPs, by default, have the same preference level. To prefer one LSP over another, set different preference levels for individual LSPs. The LSP with the lowest preference value is used. The default preference of all LSPs is 7, which is lower (more preferred) than all learned routes except for direct interface routes.

To change the default preference value, include the preference statement at the [edit protocols mpls label-switched-path *lsp-path-name*] or [edit protocols mpls label-switched-path *lsp-path-name* (primary | secondary)] hierarchy level:

```
preference preference;
```

Configure Whether to Record Path Routes

The JUNOS implementation of RSVP supports the Record Route Object, which allows an LSP to actively record the routers through which it transits. You can use this information for diagnostic purposes and to prevent routing loops. By default, path route information is recorded. To disable recording, include the no-record statement within the label-switched-path statement at the [edit protocols mpls label-switched-path *lsp-path-name*] or [edit protocols mpls label-switched-path *lsp-path-name* (primary | secondary)] hierarchy level.

```
no-record;
```

Configure the MPLS CoS Value

When IP traffic enters an LSP tunnel, the ingress router marks all packets with a class-of-service (CoS) value, which is used to place the traffic into a transmission priority queue. On the router, for SDH/SONET and T3 interfaces, each interface has four transmit queues. The CoS value is encoded as part of MPLS header and remains in the packets until the MPLS header is removed when the packets exit from the egress router. The routers within the LSP utilize the CoS value set at the ingress router.

MPLS class of service works in conjunction with the router's general CoS functionality. If you do not configure any CoS features, the default general CoS settings are used. For MPLS class of service, you might want to prioritize how the transmit queues are serviced by configuring weighted round-robin and to configure congestion avoidance using Random Early Detection (RED). The general CoS features are described in the *JUNOS Internet Software Configuration Guide: Interfaces and Chassis*.

When traffic enters an LSP tunnel, the CoS bits in the MPLS header are set in one of two ways. In the first way, the number of the output queue into which the packet was buffered is written into the MPLS header and is used as the packet's CoS value. This behavior is the default, and no configuration is required. The *JUNOS Internet Software Configuration Guide: Interfaces and Chassis* explains the IP CoS values, and summarizes how the CoS bits are treated.

In the second way, you set a fixed CoS value on all packets entering the LSP tunnel. This means that all packets entering the LSP receive the same class of service. To do this, include the class-of-service statement at the [edit protocols mpls label-switched-path *lsp-path-name*] or [edit protocols mpls label-switched-path *lsp-path-name* (primary | secondary)] hierarchy level

```
class-of-service cos-value;
```

The CoS value can be a decimal number from 0 through 7. This number corresponds to a three-bit binary number. The high-order two bits of the CoS value select which transmit queue to use on the outbound interface card.

The low-order bit of the CoS value is treated as the packet loss priority (PLP) bit and is used to select the RED drop profile to use on the output queue. If the low-order bit is 0, the non-PLP drop profile is used, and if the low-order bit is 1, the PLP drop profile is used. It is generally expected that RED more aggressively drops packets that have the PLP bit set. For more information about RED and drop profiles, see the *JUNOS Internet Software Configuration Guide: Interfaces and Chassis*.



Note

If you configure the PLP drop profile to drop packets more aggressively, then setting the CoS value to 7, for example, causes traffic to have a lesser chance of getting through than if you set the value to 6. Keep this in mind when configuring drop profiles.

Table 1 summarizes how MPLS CoS values correspond to the transmit queue and PLP bit. Note that in MPLS, the mapping between the CoS bit value and the output queue is hard-coded. You cannot configure the mapping for MPLS; you can configure it only for IPv4 traffic flows, as described in the *JUNOS Internet Software Configuration Guide: Interfaces and Chassis*.

Table 1: MPLS CoS Values

MPLS CoS Value	Bits	Transmit Queue	PLP Bit
0	000	0	Not set
1	001	0	Set
2	010	1	Not set
3	011	1	Set
4	100	2	Not set
5	101	2	Set
6	110	3	Not set
7	111	3	Set

Because the CoS value is part of the MPLS header, the value is associated with the packets only as they travel through the LSP tunnel. The value is not copied back to the IP header when the packets exit from the LSP tunnel.

Configure an LSP to be Adaptive

An LSP occasionally might need to reroute itself. Reasons include the following:

- Continuous reoptimization process is configured with the `optimize-timer` statement.

- The current path has connectivity problems.

- The LSP is preempted by another LSP, as configured with the `priority` statement, and is forced to reroute.

- The explicit-path information for an active LSP is modified, or the LSP's bandwidth is increased.

You can configure an LSP to be *adaptive* when it is attempting to reroute itself. When it is adaptive, the LSP holds onto existing resources until the new path is successfully established and traffic has been cut over to the new LSP. To retain its resources, an adaptive LSP does the following:

- Maintains existing paths and allocated bandwidths—This ensures that the existing path is not torn down prematurely and allows the current traffic to continue flowing while the new path is being set up.

- Avoids double counting for links that share the new and old paths—Double-counting occurs when an intermediate router does not recognize that the new and old paths belong to the same LSP and counts them as two separate LSPs, requiring separate bandwidth allocations. If some links are close to saturation, double-counting might cause the setup of the new path to fail.

By default, adaptive behavior is disabled. You can include the `adaptive` statement in two different hierarchy levels. If you specify the `adaptive` statement at the LSP hierarchy level [`edit protocols mpls label-switched-path lsp-path-name`], then adaptive will be effective on all primary/secondary paths of the LSP. This means both the primary and secondary paths share the same bandwidth on common links.

```
[edit protocols mpls label-switched-path lsp-path-name]
adaptive;
```

If you specify the adaptive statement at the primary/secondary hierarchy level [edit protocols mpls label-switched-path *lsp-path-name* (primary | secondary)], then adaptive will be effective only on the path it is specified. Bandwidth double counting will happen between different paths.

```
[edit protocols mpls label-switched-path lsp-path-name (primary | secondary)]
adaptive;
```

Configure Priority and Preemption

When there is insufficient bandwidth to establish a more important LSP, you might want to tear down a less important existing LSP to free up the bandwidth. You do this by preempting the existing LSP.

Whether an LSP can be preempted is determined by two properties associated with the LSP:

Setup priority—Determines whether a new LSP that preempts an existing LSP can be established. For preemption to occur, the setup priority of the new LSP must be higher than that of the existing LSP. Also, the act of preempting the existing LSP must produce sufficient bandwidth to support the new LSP. That is, preemption occurs only if the new LSP can be set up successfully.

Hold priority—Determines the degree to which an LSP holds onto its session reservation after the LSP has been set up successfully. When the hold priority is high, the existing LSP is less likely to give up its reservation and hence it is unlikely that the LSP can be preempted.

You cannot configure an LSP with a high setup priority and a low hold priority because permanent preemption loops might result if two LSPs are allowed to preempt each other. You must configure the hold priority to be higher than or equal to the setup priority.

The setup priority also defines the relative importance of LSPs on the same ingress router. When the software starts, when a new LSP is established, or during fault recovery, the setup priority determines the order in which LSPs are serviced. Higher priority LSPs tend to be established first and hence enjoy more optimal path selection.

To configure the LSP's preemption properties, include the priority statement at the [edit protocols mpls label-switched-path *lsp-path-name*] or [edit protocols mpls label-switched-path *lsp-path-name* (primary | secondary)] hierarchy level:

```
priority setup-priority hold-priority;
```

Both *setup-priority* and *hold-priority* can be a value from 0 through 7. The value 0 corresponds to the highest priority, and the value 7 to the lowest. By default, an LSP has a setup priority of 7 (that is, it cannot preempt any other LSPs) and a hold priority of 0 (that is, other LSPs cannot preempt it). These defaults are such that preemption does not happen. When you are configuring these values, the setup priority should always be less than or equal to the hold priority.

Optimize Signaled LSPs

Once an LSP has been established, topology or resources changes might, over time, make the path suboptimal. A subsequent recomputation might be able to determine a more optimal path.

If reoptimization is enabled, an LSP can be rerouted through different paths by constrained-path recomputations. However, if reoptimization is disabled, the LSP has a fixed path and cannot take advantage of newly available network resources. The LSP is fixed until the next topology change breaks the LSP and forces a recomputation.

Reoptimization is not related to fail-over. A new path is always computed when topology failures occur that disrupt an established path.

Because of the potential system overhead involved, you need to control carefully the frequency of reoptimization. Network stability might suffer when reoptimization is enabled. By default, `optimize-timer` is set to 0 (that is, it is disabled).

Configuring LSP optimization is meaningful only when constrained-path LSP computation is enabled, which is the default behavior. For more information about constrained-path LSP computation, see “Disable Constrained Path LSP Computation” on page 52.

To enable path reoptimization, include the `optimize-timer` statement at the [edit protocols mpls label-switched-path *lsp-path-name*] or [edit protocols mpls label-switched-path *lsp-path-name* (primary | secondary)] hierarchy level:

```
optimize-timer seconds;
```

After reoptimization is run, the result is accepted only if it meets the following criteria:

1. The new path is not higher in IGP metric. (The metric for the old path is updated during computation, so if a recent link metric changed somewhere along the old path, it is accounted for).
2. If the new path has the same IGP metric, it is not more hops away.
3. The new path does not cause preemption. (This is to reduce the ripple effect of preemption causing more preemption).
4. The new path does not worsen congestion overall. This is done by comparing the percentage of available bandwidth on each link traversed by the new and old paths, starting from the most congested links.

While all the above excluding conditions are met, then:

5. If the new path has a lower IGP metric, accept it.
6. If the new path has an equal IGP metric and lower hop count, accept it.
7. If you choose least-fill as a load-balancing algorithm and if the new path reduces congestion by at least 10 percent aggregated over all links it traversed, then accept it. For random or most-fill algorithms, this rule does not apply.
8. Otherwise, reject the new path.

To disable items 2, 3, 4 and 6 above, issue the `clear mpls optimize-aggressive` statement or under `[config protocols mpls]`, configure the `optimize-aggressive` statement:

```
optimize-aggressive;
```

Including the `optimize-aggressive` statement makes the reoptimization process more aggressive. Not only does it tend to reroute more often, it also limits the reoptimization algorithm to be based on the IGP metric only.

Configure the Maximum Path Length

By default, each LSP can traverse a maximum of 255 hops, including the ingress and egress routers. To modify this value, include the `hop-limit` statement at the `[edit protocols mpls label-switched-path lsp-path-name]` or `[edit protocols mpls label-switched-path lsp-path-name (primary | secondary)]` hierarchy level:

```
hop-limit number;
```

The number of hops can be from 2 through 255. (A path with two hops consists of the ingress and egress routers only.)

Configure the Path Bandwidth

Each LSP has a bandwidth value. This value is included in the sender's `Tspec` field in RSVP path setup messages. To specify a bandwidth value, include the `bandwidth` statement at the `[edit protocols mpls label-switched-path lsp-path-name]` or `[edit protocols mpls label-switched-path lsp-path-name (primary | secondary)]` hierarchy level.

```
bandwidth bps;
```

You specify the bandwidth value in bits per second, with a higher value implying a greater user traffic volume. The default bandwidth is 0 bits per second.

A nonzero bandwidth requires transit routers to reserve capacity along the outbound links for the path. This is done using RSVP's reservation scheme. Any failure in bandwidth reservation (such as failures at RSVP policy control or admission control) might cause the LSP setup to fail.

Configure the Standby State

By default, secondary paths are set up only as needed. To have the system maintain a secondary path in a hot-standby state indefinitely, include the `standby` statement at the `[edit protocols mpls label-switched-path lsp-path-name secondary]` hierarchy level:

```
[edit protocols mpls label-switched-path lsp-path-name secondary]
standby;
```

The hot-standby state is meaningful only on secondary paths.

Maintaining a path in a hot-standby state enables swift cutover to the secondary path when downstream routers on the current active path indicate connectivity problems.

The hot-standby state has two advantages:

It eliminates the call-setup delay during network topology changes. Call setup can suffer from significant delays when network failures trigger large numbers of LSP reroutes at the same time.

A cutover to the secondary path can be made before RSVP learns that an LSP is down. There can be significant delays between the time the first failure is detected by protocol machinery (which can be an interface down, a neighbor's becoming unreachable, a route's becoming unreachable, or a transient routing loop's being detected) and the time an LSP actually fails (which requires a timeout of soft state information between adjacent RSVP routers). When topology failures occur, hot-standby secondary paths can usually achieve the smallest cutover delays with minimum disruptions to user traffic.

When the primary path is again considered to be stable, traffic is automatically switched from the standby secondary path back to the primary path. The switch is performed no faster than twice the retry-timer interval and only if the primary path exhibits stability throughout the entire switch interval.

The drawback of the hot-standby state is that more state information must be maintained by all the routers along the path, which requires overhead from each of the routers.

Configure LSP Hold Time

When an LSP transitions from being up to being down, or from down to up, this transition takes effect immediately in the router software and hardware. However, when advertising LSPs into IS-IS, you may want to damp LSP transitions, thereby not advertising the transition until a certain period of time has transpired (known as the hold time). In this case, if the LSP goes from up to down, the LSP is not advertised as being down until it has remained down for the hold-time period. Transitions from down to up are advertised into IS-IS immediately. Note that LSP damping only affects IS-IS advertisements of the LSP; other routing software and hardware react immediately to LSP transitions.

To damp LSP transitions, you can include the advertisement-hold-time statement at the [edit protocols mpls] hierarchy level:

```
[edit protocols mpls]
  advertise-hold-time seconds;
```

The *seconds* can be a value from 0 to 65535 seconds. The default is 5 seconds.

Configure LDP Tunneling

To correctly identify an LDP session associated with an RSVP LSP, ensure that the RSVP LSP endpoint address is the same as the transport address of the LDP peer.

Configure Fate Sharing

Fate sharing allows you to create a database of information that CSPF uses to compute one or more backup paths to use in case the primary path becomes unstable. The database describes the relationships between elements of the network, such as routers and links.

Through fate sharing, you can now configure backup paths that minimize the number of shared links and fiber paths with the primary paths as much as possible, to ensure that in the event of a fiber cut, the minimum amount of data is lost and that a path still exists to the destination.

For a backup path to work optimally, it must not share links or physical fiber paths with the primary path. This ensures that a single point of failure will not affect the primary and backup paths at the same time.

To configure fate sharing, include the fate-sharing statement at the [edit routing-options] hierarchy level:

```
[edit routing-options]
fate-sharing {
  group group-name;
  cost value;
  from address <to address>;
}
```

Each fate-sharing group must have a name, which can be up to 32 characters long and can contain letters, digits, periods (.) and hyphens (-). You can define up to 512 groups.

Fate-sharing groups contain three types of objects:

Point-to-point links—Identified by the IP addresses at each end of the link. Unnumbered point-to-point links are typically identified by borrowing IP addresses from other interfaces. Order is not important; from 1.2.3.4 to 1.2.3.5 and from 1.2.3.5 to 1.2.3.4 have the same meaning.

Non-point-to-point links—Include links on a LAN interface (such as Gigabit Ethernet interfaces), or NBMA interfaces, (such as ATM or Frame Relay). You identify these links by their individual interface address. For example, if a LAN 192.168.200.0/24 has four routers attached to it, each router link is individually identified:

```
from 192.168.200.1; # LAN interface of router 1
from 192.168.200.2; # LAN interface of router 2
from 192.168.200.3; # LAN interface of router 3
from 192.168.200.4; # LAN interface of router 4
```

Sequence is insignificant; you can list the addresses in any order.

A router node—Identified by its configured router ID.

All objects in a group share certain similarities. For example, you can define a group for all fibers sharing the same fiber conduit, all optical channels that share the same fiber, all links that connect to the same LAN switch, all equipment sharing the same power source, and so on. All objects are treated as /32 host addresses.

For a group to be meaningful, it should contain at least two objects. You can configure groups with zero or one object; these groups are ignored during processing.

An object can be in any number of groups, and a group can contain any number of objects. Each group has a configurable cost attributed to it, which represents the level of impact this group has on CSPF computations. The higher the cost, the less likely a backup path will share with the primary path any objects in the group. The cost is directly comparable to traffic engineering metrics. By default, the cost is 1. Changing the fate-sharing database does not affect existing established LSPs until the next re-optimization of CSPF.

Implications to CSPF

When CSPF computes the primary paths of an LSP (or secondary paths when the primary path is not active), it ignores the fate-sharing information. You always want to find the best possible path (least IGP cost) for the primary path.

When CSPF computes a secondary path while the primary path (of the same LSP) is active, the following occurs:

1. CSPF identifies all fate-sharing groups that are associated with the primary path. CSPF does this by identifying all links and nodes that the primary path traverses through and compiling group lists that contain at least one of the links or nodes. CSPF ignores the ingress and egress nodes in the search.
2. CSPF checks each link in the TED against the compiled group list. If the link is a member of a group, the cost of the link is increased by the cost of the group. If a link is a member of multiple groups, then all group costs are added together.
3. CSPF performs the check for every node in the TED, except the ingress and egress node. Again, a node can belong to multiple groups, so costs are additive.
4. The router performs regular CSPF computation with the adjusted topology.

Example: Configure Fate Sharing

Configure fate-sharing groups thunder and shadow. Because shadow has no objects, it is ignored during processing.

```
[edit routing-options]
  fate-sharing {
    group thunder {
      cost <20>;                # optional, default value is 1
      from 1.2.3.4 to 1.2.3.5;  # a point-to-point link
      from 192.168.200.1;      # LAN interface
      from 192.168.200.2;      # LAN interface
      from 192.168.200.3;      # LAN interface
      from 192.168.200.4;      # LAN interface
      from 10.168.1.220;        # Router ID of a router node
      from 10.168.1.221;        # Router ID of a router node
    }
    group shadow {
    }
  }
}
```

Configure All Other MPLS Routers for Signaled LSPs

To configure signaled LSPs on all the MPLS routers that should participate in MPLS, you need to enable MPLS and RSVP on these routers, as described in “Minimum MPLS Configuration” on page 37 and “Enable RSVP” on page 64.

Enable RSVP

For all routers that you want to have participate in signaled LSPs, you must enable RSVP because it is used to set up LSPs. To do this, include the following statements in the configuration. In general, we recommend that you enable RSVP on all router interfaces, except for those on the AS border:

```
[edit]
interfaces {
  interface-name {
    unit logical-unit-number {
      family mpls;
    }
  }
}
protocols {
  mpls {
    interface all;
  }
  rsvp {
    interface all;
  }
}
```

For more information about RSVP, see “RSVP Configuration Guidelines” on page 117.

Examples: Configure Signaled LSPs

On the ingress router, create a constrained path LSP in which the JUNOS software makes all the forwarding decisions. When the LSP is successfully set up, a route toward 11.1.1.1/32 is installed in the inet.3 table so that all BGP routes with matching BGP next-hop addresses can be forwarded through the LSP.

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  rsvp {
    interface so-0/0/0;
  }
  mpls {
    label-switched-path to-hastings {
      to 11.1.1.1;
    }
    interface so-0/0/0;
  }
}
```

On the ingress router, create an explicit-path LSP and specify the transit routers between the ingress and egress routers. In this configuration, no constrained-path computation is performed. For the primary path, all intermediate hops are strictly specified so its route cannot change. The secondary path must travel through router 14.1.1.1 first, then take whatever route is available to reach the destination. The remaining route taken by secondary path is typically the shortest path computed by the IGP.

```
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family mpls;
    }
  }
}
protocols {
  rsvp {
    interface so-0/0/0;
  }
  mpls {
    path to-hastings {
      14.1.1.1 strict;
      13.1.1.1 strict;
      12.1.1.1 strict;
      11.1.1.1 strict;
    }
    path alt-hastings {
      14.1.1.1 strict;
      11.1.1.1 loose; # Any IGP route is acceptable
    }
    label-switched-path hastings {
      to 11.1.1.1;
      hop-limit 32;
      bandwidth 10m; # Reserve 10 mbps
      no-cspf; # do not perform constrained-path computation
      primary to-hastings;
      secondary alt-hastings;
    }
  }
  interface so-0/0/0;
}
}
```

On the ingress router, create a constrained-path LSP in which the JUNOS software makes most of the forwarding decisions, taking into account the hop constraints listed in the path statements. The LSP is adaptive so that no double bandwidth counting occurs on links shared by primary and secondary paths. To acquire the necessary link bandwidth, this LSP is allowed to preempt lower priority sessions. Finally, this path always keeps the secondary path in hot-standby state for quick failover.

```
[edit protocols]
mpls {
  path to-hastings {
    14.1.1.1 loose;
  }
  path alt-hastings {
    12.1.1.1 loose;
    11.1.1.1 strict;
  }
  label-switched-path hastings {
    to 11.1.1.1;
    bandwidth 10m;      # Reserve 10 mbps
    priority 0 0;      # Preemptive, but not preemptable
    adaptive;          # Set adaptivity
    primary to-hastings;
    secondary alt-hastings{
      standby;
      bandwidth 1m;    # Reserve only 1 Mbps for the secondary path
    }
  }
}
interface all;
}
```

On the ingress router, create a constrained-path LSP in which the JUNOS software makes most the forwarding decisions for the primary path, subject to constraints of the path to-hastings, and in which the secondary path is an explicit path. The primary path must transit green or yellow links and must stay away from red links. The primary path is periodically recomputed and reoptimized. Finally, this path always keeps the secondary path in hot-standby state for quick failover.

When the LSP is up—either because the primary or secondary path is up, or both are up—the prefix 16.0.0.0/8 is installed in the inet.3 table so that all BGP routes whose BGP next hop falls within that range can use the LSP. Also the prefix 17/8 is installed in the inet.0 table so that BGP can only resolve its next hop through it and the route also can be reached using traceroute or ping. These two routes are in addition to the 11.1.1.1/32 route.

```
[edit protocols]
mpls {
  admin-groups {
    green 1;
    yellow 2;
    red 3;
  }
  path to-hastings {
    14.1.1.1 loose;
  }
  path alt-hastings {
    14.1.1.1 strict;
    13.1.1.1 strict;
    12.1.1.1 strict;
    11.1.1.1 strict;
  }
  label-switched-path hastings {
    to 11.1.1.1;
    bandwidth 100m;
    install 16.0.0.0/8;      # in inet.3; cannot use to traceroute or ping
    install 17.0.0.0/8 active; # installed in inet.0; can use to traceroute or ping
    primary to-hastings {
      admin-group {          # further constraints for path computation
        include [ green yellow ];
        exclude red;
      }
      optimize-timer 3600; # reoptimize every hour
    }
    secondary alt-hastings {
      standby;
      no-cspf;              # do not perform constrained-path computation
    }
  }
  interface all;
}
```

