AT&T

MPLS Configuration Guide, 17.2.0

November 2017

Supporting AT&T Vyatta Network Operating System

Contents

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Copyright Statement

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property. All other marks are the property of their respective owners.

The training materials and other content provided herein for assistance in training on the Vyatta vRouter may have references to Brocade as the Vyatta vRouter was formerly a Brocade product prior to AT&T’s acquisition of Vyatta. Brocade remains a separate company and is not affiliated to AT&T.

MPLS Overview

Multiprotocol Label Switching (MPLS) on the AT&T Vyatta vRouter provides the data plane support and the basic Label Distribution Protocol (LDP) and Resource Reservation Protocol for Traffic Engineering (RSVP-TE) support that serves as a foundation on which to build additional features.

The following support is provided:

- Basic MPLS forwarding
- LDP
- RSVP-TE

Basic MPLS forwarding support

The following basic MPLS forwarding features are supported:

- Label Switched Paths (LSPs) using IPv4 next hops
- Imposition of labels onto IPv4 packets
- Routes with de-aggregate to global tables
- ECMP for imposition and labeled routes
- IP fragmentation
- IP traceroute using LSPs (when TTL propagation is enabled)
- Generation of too-big ICMP (when the MTU is exceeded with the DF bit set)
- Configurable TTL propagation
- Configurable label table size
- Per-interface packet statistics

LDP support

The following LDP features are supported:

- Label allocation and advertisement control
  - Outbound filtering
  - Allocation control; only label ranges supported
- Null label allocation
  - Explicit null if configured; implicit null if not configured
  - Session authentication
- IGP synchronization
- Downstream-on-demand advertisement mode
- IPv4 LSPs
- ECMP for imposition and labeled routes
- Graceful restart helper
- Operation in default VRF only

RSVP-TE support

The following RSVP features are supported:

- CSPF OSPF computation (resource utilization)
- Explicit path (loose and strict)
- Fast reroute with one-to-one protection
- IGP shortcut (IGP routes through a tunnel)

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

- IPv4 point-to-point tunnels
- Node and link protection
- Path affinities
- Reoptimization
- Record route
- Refresh reduction (scalability)
- Shared explicit filter
- Tail end implicit or explicit null signaling
- RSVP operation in default VRF only
- Graceful restart
- Tunnels to destinations outside the OSPF area

Not yet supported

The following features are not currently supported:

- L3VPN
- Diffserv-aware TE
- Facility backup
- MPLS QoS
- Point-to-multipoint tunnels
- Shared Risk Link Groups (SRLG)

IETF RFC and Internet Draft support for MPLS

The AT&T Vyatta vRouter implementation of MPLS supports the following IETF RFCs and Internet drafts.

RFC Number	RFC Title
RFC 3031	Multiprotocol Label Switching Architecture
RFC 3032	MPLS Label Stack Encoding
RFC 5036	LDP Specification
RFC 2205	Resource ReSerVation Protocol (RSVP) Version 1 Functional Specification
RFC 2209	Resource ReSerVation Protocol (RSVP) Version 1 Message Processing Rule
RFC 3209	RSVP-TE
RFC 4090	Facility backup and Fast Reroute
RFC 3630	TE Extensions to OSPF v2

How MPLS works

MPLS uses a label-switching forwarding method to direct packets through a network. In label switching, a packet is assigned a label and passes along a predetermined path of routers. Forwarding decisions are based on the contents of the label, rather than information in the IP header of the packet.

The following sections describe these basic MPLS concepts:

- How packets are forwarded through an MPLS domain
- The kinds of Label Switched Paths (LSPs) that can be configured on a device
- The components of an MPLS label header

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

How packets are forwarded through an MPLS domain

In an MPLS domain, packets are forwarded from one MPLS-enabled router to another along a predetermined path, called an LSP.

An MPLS domain consists of a group of MPLS-enabled routers, called Label Switching Routers (LSRs). LSPs are one-way paths between MPLS-enabled routers on a network. To provide two-way traffic, you must configure LSPs in each direction.

The LSRs at the head end and tail end of an LSP are known as Label Edge Routers (LERs). The LER at the head end, where packets enter the LSP, is known as the ingress LER . The LER at the tail end, where packets exit the LSP, is known as the egress LER. Each LSP has one ingress LER and one egress LER. Packets in an LSP flow in one direction: from the ingress LER toward the egress LER. Between the ingress and egress LERs, there may be zero or more transit LSRs. A device that is enabled for MPLS can perform the role of an ingress LER, transit LSR, or egress LER in an LSP. Furthermore, a device can serve simultaneously as an ingress LER for one LSP, transit LSR for another LSP, and an egress LER for some other LSP.

The figure titled "Label switching in an MPLS domain" depicts an MPLS domain with a single LSP consisting of three LSRs: an ingress LER, a transit LSR, and an egress LER.

Note: In the following figure, Ingress LER, Transit LSR, and Egress LER are virtual routers.

Figure 1: Label switching in an MPLS domain

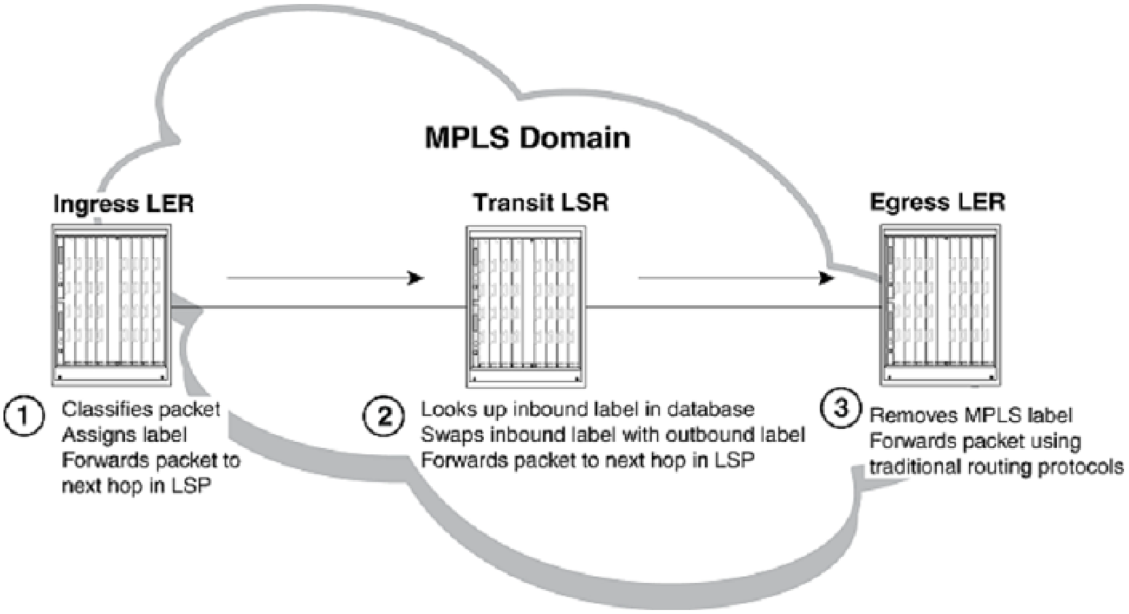

Label switching in an MPLS domain works as described below.

1. The Ingress LER receives a packet and pushes a label onto it.
   When a packet arrives on an MPLS-enabled interface, the device determines to which LSP (if any) the packet is assigned. Specifically, the device determines to which Forwarding Equivalence Class (FEC) the packet belongs. An FEC is simply a group of packets that are all forwarded in the same way, typically represented by a prefix in the routing table. FECs are mapped to LSPs. When a packet belongs to an FEC and an LSP is mapped to that FEC, the packet is assigned to the LSP.

   When a packet is assigned to an LSP, the device, acting as an ingress LER, applies (pushes) a tunnel label onto the packet. A label is a 32-bit, fixed-length identifier that is significant only to MPLS. Refer to the MPLS label header encoding information for specific information about the contents of a label. From this point until the packet reaches the egress LER at the end of the path, the packet is forwarded by using information in its label, not information in its IP header. The IP header of the packet is not examined again as long as the packet traverses the LSP.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

On the ingress LER, the label is associated with an outbound interface. After receiving a label, the packet is forwarded over the outbound interface to the next router in the LSP.

2. A transit LSR receives the labeled packet, swaps the label, and forwards the packet to the next LSR.
   In an LSP, zero or more transit LSRs can exist between the ingress and egress LERs. A transit LSR swaps labels on an MPLS packet and forwards the packet to the next router in the LSP.

   When a transit LSR receives an MPLS packet, it looks up the label in its MPLS label table. This table maps the label and inbound interface to a new label and outbound interface. The transit LSR replaces the old label with the new label and sends the packet out the outbound interface that is specified in the table. This process repeats at each transit LSR until the packet reaches the next-to-last LSR in the LSP (for signaled LSPs).

   The following figure illustrates an example of the label-swapping process on a transit LSR.

   Figure 2: Label swapping on a transit LSR

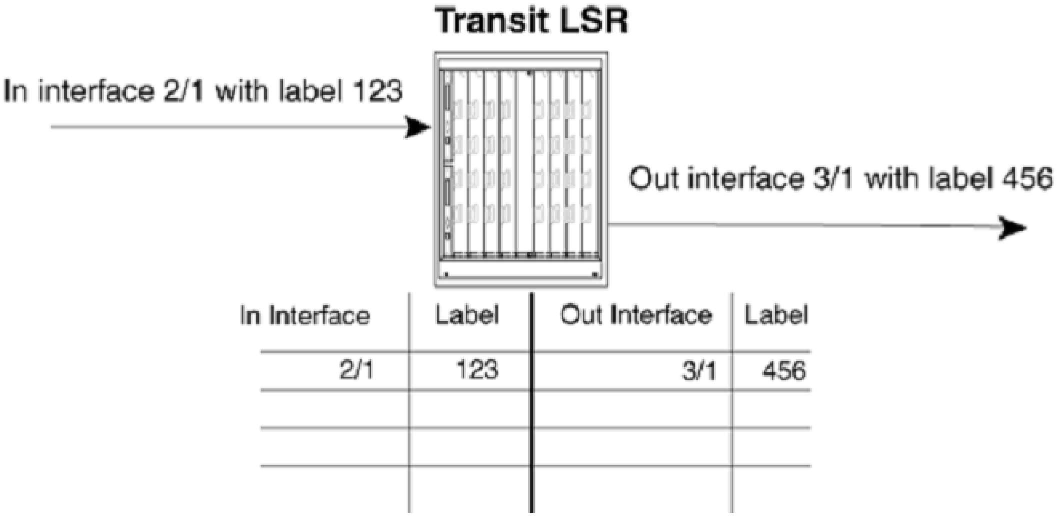
Transit LSR
In interface 2/1 with label 123
Out interface 3/1 with label 456

In Interface	Label	Out Interface	Label
2/1	123	3/1	456

   In this example, a packet comes into interface 2/1 with label 123. The transit LSR then looks up this interface-label pair in its MPLS label table. The inbound interface-label pair maps to an outbound interface-label pair - in this example, interface 3/1 with label 456. The LSR swaps label 123 with label 456 and forwards the packet out interface 3/1.

3. The egress LER receives the labeled packet, pops the label, and forwards the IP packet.
   When the packet reaches the egress LER, the MPLS label is removed (this is called popping the label), and the packet can then be forwarded to its destination by using standard hop-by-hop routing protocols. On signaled LSPs, the label is popped at the penultimate (next to last) LSR, rather than the egress LER.

Penultimate hop popping

On signaled LSPs, the MPLS label is popped at the next-to-last LSR in the LSP, instead of at the egress LER. This action is called penultimate hop popping.

Penultimate hop popping improves forwarding efficiency by allowing the egress LER to avoid performing both an MPLS forwarding table lookup and an IP forwarding table lookup for each packet exiting the LSP. Instead, the MPLS label is popped at the penultimate LSR, and the packet is forwarded to the egress LER with no MPLS encoding. The egress LER, in fact, does not recognize the packet as emerging from an LSP.

The following figure illustrates the operation that takes place at the penultimate LSR in an LSP.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Figure 3: Penultimate hop popping

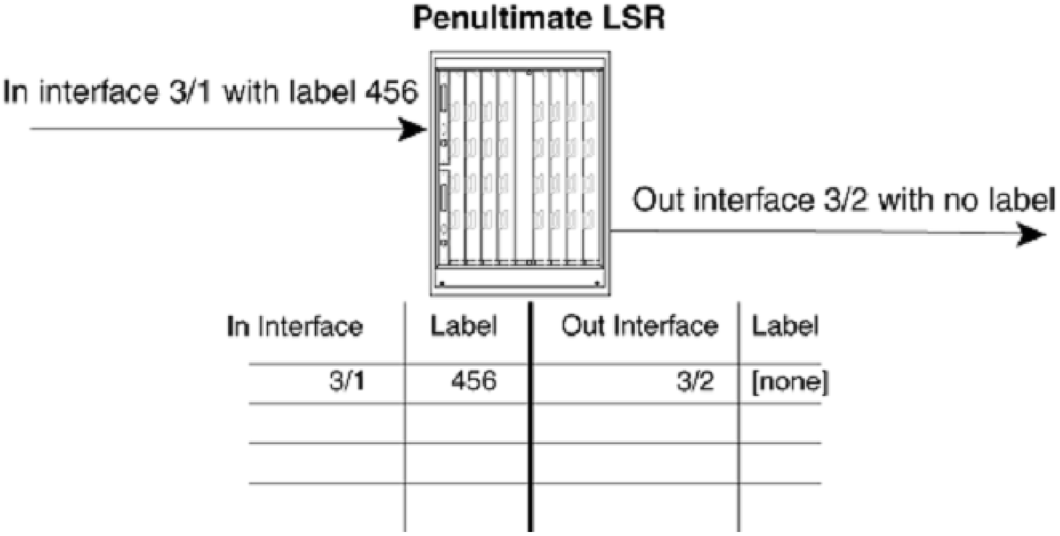
Penultimate LSR
In interface 3/1 with label 456
Out interface 3/2 with no label

In Interface	Label	Out Interface	Label
3/1	456	3/2	[none]

When an LSR receives an MPLS packet, it looks up the label in its MPLS forwarding table. Normally, this table maps the label and inbound interface to a new label and outbound interface. However, when this LSR is the penultimate LSR in an LSP, the label and inbound interface map only to an outbound interface. The penultimate LSR pops the label and forwards the packet, now a regular IP packet, out the outbound interface. When the packet reaches the egress LER, no indication exists that it was forwarded over an LSP. The packet is forwarded by using standard hop-by-hop routing protocols.

MPLS label header encoding

The following diagram illustrates the structure of the 32-bit MPLS label header. When a packet enters an LSP, the ingress LER pushes a label onto the packet.

Figure 4: Structure of an MPLS Label Header

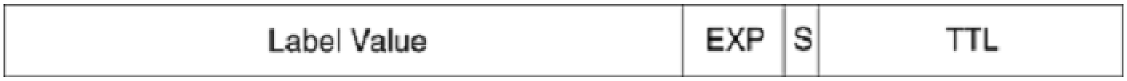

Label Value	EXP	S	TTL

An MPLS label header comprises of the following parts:

Label value (20 bits)

The label value is an integer in the range of 16 through 1048575. (Labels 0 through 15 are reserved by the IETF for special usage.)

EXP field (3 bits)

The EXP field is designated for experimental usage. By default, a device uses the EXP field to define a Class of Service (CoS) value for prioritizing packets traveling through an LSP.

S (Bottom of Stack) field (one bit)

An MPLS packet can be assigned multiple labels. When an MPLS packet has multiple labels, they are logically organized in a last-in, first-out label stack. An LSR performs a pop or swap operation on the topmost label; that is, the most recently applied label in the stack. The Bottom of Stack field indicates whether this label is the last (oldest) label in the stack. When the label is the last one in the stack, the Bottom of Stack field is set to one. If not, the Bottom of Stack field is set to zero.

A device acting as an LSR can perform one push, swap, or pop operation on an incoming MPLS packet. The device can accept MPLS packets that contain multiple labels, but only the topmost label is acted upon.

TTL field (eight bits)

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

The TTL field indicates the Time To Live (TTL) value for the MPLS packet. At the ingress LER, the TTL value of an IP packet is copied to its MPLS TTL field. At each transit LSR hop, the MPLS TTL value is decremented by one. When the MPLS TTL value reaches zero, the packet is discarded. The MPLS TTL value is copied into the IP header at the egress LER.

Displaying Global MPLS Information

You can display the following information about the global MPLS configuration:

- Information about MPLS-enabled interfaces on the device
- Statistics about the MPLS-enabled interfaces
- MPLS summary information
- Status information about signaled LSPs that are configured on the device
- Information about paths that are configured on the device
- The label that is applied at each hop in an LSP
- Contents of the MPLS routing table

Displaying the label assignment for MPLS traffic leaving the LSR

To display the MPLS out segment table that shows the label assignment for MPLS traffic leaving the LSR, enter the `show mpls out-segment-table` command.

```
vyatta@vyatta:~$ show mpls out-segment-table
      Out-segment with ix: 12, owner: LDP, out intf: dp0p1s1, out label: 3
    Nexthop addr: 192.166.3.2        cross connect ix: 1, op code: Push

       Out-segment with ix: 4, owner: LDP, out intf: dp0s4, out label: N/A
    Nexthop addr: 192.168.252.253        cross connect ix: 3, op code: Push
```

Displaying the MPLS forwarding information

To display MPLS forwarding information, enter the `show mpls forwarding-table` command.

```
vyatta@vyatta:~$ show mpls forwarding-table
  Codes: > - selected FTN, p - stale FTN, B - BGP FTN, K - CLI FTN
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN,
       I - IGP-Shortcut, U - unknown FTN
   FEC                        Nexthop               Out-Label Out-Intf
>L 1.1.1.1/32                 192.166.3.2           52490     dp0p1s1
                              192.166.1.1           3         dp0p1s3
>L 3.3.3.3/32                 192.166.3.2           3         dp0p1s1
>L 192.166.2.0/30             192.166.3.2           3         dp0p1s1
>L 192.166.4.0/30             192.166.3.2           3         dp0p1s1
>L 192.166.8.0/30             192.166.3.2           3         dp0p1s1
```

Displaying the MPLS cross-connect table

Use the `show mpls cross-connect` command to show the association between the MPLS in segment table and the MPLS out segment table.

To display the MPLS cross-connect table, enter the `show mpls cross-connect` command.

```
vyatta@vyatta:~$ show mpls cross-connect
     Cross connect ix: 3, in intf: dp0p1s1 in label: 52484 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, out intf: dp0s4, out label: N/A
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
    Nexthop addr: 192.168.252.253        cross connect ix: 3, op code: Pop
```

Displaying the MPLS label table

Use the `show mpls label-table` command to display the MPLS label table that is used for forwarding encapsulated traffic.

To display the MPLS label table, enter the `show mpls label-table` command.

```
vyatta@vyatta:~$ show mpls label-table
Codes: > - selected ILM, p - stale ILM, K - CLI ILM, T - MPLS-TP

Code FEC                 In-Lbl    Out-Lbl    In-Intf          Out-Intf         Nexthop
 >   3.3.3.3/32          52481     3          dp0p1s3          dp0p1s1          192.166.3.2
 >   10.0.0.0/8          52484     N/A        dp0p1s1          dp0s4
 192.168.252.253
 >   192.166.4.0/30      52480     3          dp0p1s3          dp0p1s1          192.166.3.2
 >   10.0.0.0/8          52484     N/A        dp0p1s3          dp0s4
 192.168.252.253
 >   192.166.8.0/30      52482     3          dp0p1s3          dp0p1s1          192.166.3.2
 >   192.166.2.0/30      52485     3          dp0p1s3          dp0p1s1          192.166.3.2
```

Displaying the MPLS incoming segment table

Use the `show mpls in-segment-table` command to display the MPLS incoming label to Next Hop Label Forwarding Entry (NHLFE) mapping.

To display the MPLS incoming segment table, enter the `show mpls in-segment-table` command.

```
vyatta@vyatta:~$ show mpls in-segment-table
  Owner: LDP, # of pops: 1, fec: 3.3.3.3/32, ILM-ID: 23
     Cross connect ix: 1, in intf: dp0p1s3 in label: 52481 out-segment ix: 12
      Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Up
       Out-segment with ix: 12, owner: LDP, out intf: dp0p1s1, out label: 3
    Nexthop addr: 192.166.3.2        cross connect ix: 1, op code: Swap

  Owner: LDP, # of pops: 1, fec: 10.0.0.0/8, ILM-ID: 17
     Cross connect ix: 3, in intf: dp0p1s1 in label: 52484 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, out intf: dp0s4, out label: N/A
    Nexthop addr: 192.168.252.253        cross connect ix: 3, op code: Pop
```

Displaying the MPLS interfaces

Use the `show mpls interface` command to display the MPLS-enabled interfaces for label switching and the corresponding label range in use for local labels.

To display the MPLS-enabled interfaces, enter the `show mpls interface` command.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@vyatta:~$ show mpls interface
Interface lo
  Label switching is disabled
Interface dp0s4
  Label switching is disabled
Interface dp0p1s1
  Label switching is enabled with label-space 0
   minimum label value configured is 16
   maximum label value configured is 1048575
Interface dp0p1s2
  Label switching is enabled with label-space 0
   minimum label value configured is 16
   maximum label value configured is 1048575
Interface dp0p1s3
  Label switching is enabled with label-space 0
   minimum label value configured is 16
   maximum label value configured is 1048575

Total number of mpls interface is 3
```

Displaying MPLS interfaces on the data plane

Use the `show dataplane mpls interfaces` command to display the MPLS-enabled interfaces on the data plane.

To display the MPLS interfaces on the data plane, enter the `show mpls dataplane mpls interfaces` command.

```
vyatta@vyatta:~$ show dataplane mpls interfaces
MPLS Interfaces
dp0p1s1, ifindex: 9, mtu: 1500
        address: 192.166.3.1/30 fe80::5054:ff:fe00:201/64 192:166:3::1/64
dp0p1s2, ifindex: 10, mtu: 1500
        address: 192.166.5.1/30 192:166:5::1/64 fe80::5054:ff:fe00:202/64
dp0p1s3, ifindex: 11, mtu: 1500
        address: 192.166.1.2/30 192:166:1::2/64 fe80::5054:ff:fe00:203/64
```

Displaying the MPLS label table information on the data plane

Use the `show dataplane mpls label-table` command to display the MPLS label table information on the data plane. The with-prefix option includes FEC information in the output.

To display the MPLS label table information on the data plane, enter the `show dataplane mpls label-table` command.

```
vyatta@vyatta:~$ show dataplane mpls label-table with-prefix
Label Space: 0
in label: exp-null, fec:ipv4, outgoing label: imp-null
in label: rtr-alrt, outgoing label: imp-null
in label: exp6-null, fec:ipv6, outgoing label: imp-null
in label: 52480, fec:ipv4 192.166.4.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 52481, fec:ipv4 3.3.3.3/32
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 52482, fec:ipv4 192.166.8.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
in label: 52484, fec:ipv4 10.0.0.0/8
        nexthop via 192.168.252.253, dp0s4
in label: 52485, fec:ipv4 192.166.2.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 53120 (local), fec:ipv4 192.166.4.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 53121 (local), fec:ipv4 192.166.8.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 53128 (local), fec:ipv4 1.1.1.1/32
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: 52490
        nexthop via 192.166.1.1, dp0p1s3, outgoing label: imp-null
in label: 53129 (local), fec:ipv4 3.3.3.3/32
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 53130 (local), fec:ipv4 192.166.2.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Global MPLS Configuration Considerations

The considerations for the global configuration of MPLS are:

- TTL Propagation
- Setting global MPLS policy parameters

IP-over-MPLS TTL propagation control

In the MPLS label header, the TTL field indicates the Time To Live (TTL) value for an MPLS packet. For IP-over-MPLS applications, at the ingress LER the TTL value of an IP packet is decremented by one and the IP checksum is recalculated. The TTL value is then copied to the MPLS TTL field in the packet. At each transit LSR hop, the MPLS TTL value is decremented by one. When the MPLS TTL value reaches one or zero, the packet is discarded.

At the MPLS router that pops the label (either the penultimate LSR or egress LER), the MPLS TTL value of the incoming packet is copied to the IP TTL field of the packet, the IP TTL field is decremented by one, and the checksum is recalculated. The result is that each LSR in the MPLS domain is counted as one hop. This behavior is the default.

Optionally, you can configure TTL propagation so that the entire MPLS domain appears as two hops. In this case, the ingress LER decrements the TTL value of the IP packet by one and then places a value of 255 in the MPLS TTL field of the packet. The MPLS TTL value is decremented by one as the MPLS packet passes through each LSR in the MPLS domain. When the label is popped, the value in the MPLS TTL field is discarded, not copied to the IP TTL field of the packet. The TTL of the unlabeled IP packet is then decremented by one as it passes through the egress LER. This means that the IP TTL of the packet is decremented twice from the time it enters the ingress LER to the time it exits the egress LER, making the MPLS domain appear as two hops.

Configuring TTL propagation for the MPLS domain

By default, TTL propagation is enabled for IP over MPLS traffic when an RSVP and an LDP tunnel terminate on the same node. For traceroute to report the hops along the LSP, TTL propagation must be enabled. To make an entire MPLS domain appear as two hops, perform the following steps in configuration mode.

1. First, you must enter the following command to disable TTL propagation for IP-over-MPLS. Repeat this step for all nodes in the MPLS domain.

   Example:

   ```
   vyatta@R1# set protocols mpls disable-ip-propagate-ttl
   ```

2. Enable user-defined TTL propagation. 255 is the default. Repeat this step for all nodes in the MPLS domain.

   Example:

   Note: The `protocols mpls default-ttl` command is used only when TTL propagation is disabled.

   ```
   vyatta@R1# set protocols mpls default-ttl <0-255>
   ```

   Result: Once user-defined TTL propagation is enabled on all nodes in the MPLS domain, routes across the MPLS domain appear as 2 hops, ingress and egress, to the MPLS LSP. The `default-ttl` value is put on the packets as they enter the LSP. At the tail node, once the MPLS header is removed, the IP TTL is once again used. No details are provided on how many hops the LSP crosses in the MPLS domain.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

MPLS Commands

List of the supported MPLS operational commands and configuration commands

The supported MPLS commands are grouped into the following sections:

- MPLS operational commands
- MPLS configuration commands

MPLS operational commands

The following clear command is available with MPLS:

clear interfaces dataplane mpls counters | [<text> counters]

The following show commands are available with MPLS:

1. show mpls cross-connect
2. show mpls forwarding-table
3. show mpls in-segment-table
4. show mpls interface
5. show mpls label-table [x.x.x.x/x] [longer-prefixes]
6. show mpls out-segment-table
7. show dataplane mpls interfaces
8. show dataplane mpls label-table [with-prefix]
9. show dataplane statistics mpls

> **Note:** See also the `show mpls-ldp` and `show mpls rsvp` commands.

MPLS configuration commands

The following configuration commands are available with MPLS:

1. protocols mpls default-ttl <0-255>
2. protocols mpls disable-ip-propagate-ttl
3. protocols mpls label-range maximum-label-value <16..1048575>
4. protocols mpls label-range minimum-label-value <16..1048575>

> **Note:** See also the `protocols mpls-ldp` and `protocols mpls-rsvp` commands.

Logging is enabled for the MPLS data plane with the following command:

monitor dataplane mpls events | packet-error enable | disable

clear interfaces dataplane mpls counters

Clears the counters for MPLS-enabled interfaces on the data plane.

Syntax:

```
clear interfaces dataplane mpls [ interface ] counters
```

interface

The MPLS-enabled interface name.

Operational mode

Use this command to clear the counters for all MPLS-enabled interfaces that are displayed by the `show dataplane mpls interfaces` command.

Optionally, you can use this command to clear the counters for the named interface.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

monitor dataplane mpls

Enables or disables the generation of debug messages that are related to MPLS events or packet errors on the data plane.

Syntax:
`monitor dataplane mpls` { **`events`** | **`packet-error`** } { **`enable`** | **`disable`** }

Operational mode

Use this command to enable or disable the generation of debug messages that are related to MPLS events or packet errors on the data plane.

The following example shows how to enable MPLS data plane events monitoring.

```
vyatta@vyatta:~$ monitor dataplane mpls events enable
```

As a result, MPLS data plane event entries are logged into the data plane log, as shown with the following example of an MPLS data plane event entry in the `/var/log/dataplane/vplane.log` data plane log.

```
[  643.678900] MPLS: RTM_NEWROUTE table 254 type unicast scope 0 proto 11 in 53135 payload 4
 out 548080 dev dp0p1s12 via 10.10.12.7
```

protocols mpls default-ttl

Sets the default TTL value.

Syntax:
`set protocols mpls default-ttl` *0-255*

Syntax:
`delete protocols mpls default-ttl` *0-255*

Syntax:
`show protocols mpls default-ttl` *0-255*

0-255
The default value of the TTL, which ranges from 0 through 255.

Configuration mode

```
protocols {
    mpls {
        default-ttl <0-255>
    }
}
```

Use this command to configure the default Time To Live (TTL) value for an MPLS packet when imposing MPLS labels onto IP packets if TTL propagation is disabled. If the default TTL value is not specified, it is 255.

Note: The `protocols mpls default-ttl` command is used only when TTL propagation is disabled.

Use the `set` form of this command to specify the default TTL value.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use the `delete` form of this command to delete the previously specified default TTL value. This command resets the default TTL value to 255.

Use the `show` form of this command to display the default TTL value.

Example: Example

The following example shows how to set the TTL value to 100.

```
vyatta@R1# set protocols mpls default-ttl 100
```

protocols mpls disable-ip-propagate-ttl

Disables propagation of the TTL value from the IP packet when imposing MPLS labels.

Syntax:
`set protocols mpls disable-ip-propagate-ttl`

Syntax:
`delete protocols mpls disable-ip-propagate-ttl`

Syntax:
`show protocols mpls disable-ip-propagate-ttl`

Configuration mode

```
protocols {
    mpls {
        disable-ip-propagate-ttl
    }
}
```

Use this command to disable propagation of the Time To Live (TTL) value from the IP packet when imposing MPLS labels.

Use the `set` form of this command to disable the propagation of the TTL value from the IP packet.

Use the `delete` form of this command to re-enable the propagation of the TTL value from the IP packet.

Use the `show` form of this command to display whether the propagation of the TTL value from the IP packet is disabled.

protocols mpls label-range maximum-label-value

Sets the maximum value of the MPLS label range.

Syntax:
`set protocols mpls label-range maximum-label-value` *16..1048575*

Syntax:
`delete protocols mpls label-range maximum-label-value` *16..1048575*

Syntax:
`show protocols mpls label-range maximum-label-value` *16..1048575*

16..1048575
The maximum value of the label range, which ranges from 16 through 1048575.

Configuration mode

```
protocols {
    mpls {
        label-range {
            maximum-label-value <16..1048575>
        }
    }
}
```

The default maximum value of the label range is 1048575. The value that you set must be greater than the minimum value of the label range. Changes to the label range will take effect immediately only when no MPLS applications have been configured or if the maximum value is increased; otherwise they remain pending until the applications are disabled or the next system reboot.

Use the `set` form of this command to set the maximum value of the label range.

Use the `delete` form of this command to delete the maximum value of the label range.

Use the `show` form of this command to display the configured maximum value of the label range.

Example: Example

The following example shows how to set the maximum value of the label range to 100000.

```
vyatta@R1# set protocols mpls label-range maximum-label-value 100000
```

protocols mpls label-range minimum-label-value

Sets the minimum value of the MPLS label range.

Syntax:
`set protocols mpls label-range minimum-label-value` *16..1048575*

Syntax:
`delete protocols mpls label-range minimum-label-value` *16..1048575*

Syntax:
`show protocols mpls label-range minimum-label-value` *16..1048575*

16..1048575
The minimum value of the label range, which ranges from 16 through 1048575.

Configuration mode

```
protocols {
    mpls {
        label-range {
            minimum-label-value <16..1048575>
        }
    }
}
```

The default minimum value of the label range is 16. The value that you set must be less than the maximum value of the label range. Changes to the label range take effect immediately only when no MPLS applications have been configured or the minimum value is reduced; otherwise they remain pending until the applications are disabled or the next system reboot.

Use the `set` form of this command to set the minimum value of the label range.

Use the `delete` form of this command to delete the minimum value of the label range.

Use the `show` form of this command to display the configured minimum value of the label range.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to set the minimum value of the label range to 100.

```
vyatta@R1# set protocols mpls label-range minimum-label-value 100
```

show dataplane mpls interfaces

Displays the MPLS-enabled interfaces that are on the data plane

Syntax:

`show dataplane mpls interfaces`

Operational mode

Use this command to display the interfaces that support MPLS on the data plane.

The following example shows how to display the MPLS-enabled interfaces on the data plane.

```
vyatta@vyatta:~$ show dataplane mpls interfaces
MPLS Interfaces
dp0p1s1, ifindex: 9, mtu: 1500
        address: 192.166.3.1/30 fe80::5054:ff:fe00:201/64 192:166:3::1/64
dp0p1s2, ifindex: 10, mtu: 1500
        address: 192.166.5.1/30 192:166:5::1/64 fe80::5054:ff:fe00:202/64
dp0p1s3, ifindex: 11, mtu: 1500
        address: 192.166.1.2/30 192:166:1::2/64 fe80::5054:ff:fe00:203/64
```

show dataplane mpls label-table [with-prefix]

Displays the MPLS label table information on the data plane.

Syntax:

`show dataplane mpls label-table` [with-prefix]

with-prefix

Displays the FEC information in the output.

Operational mode

Use this command to display the label table information used for forwarding MPLS-encapsulated traffic on the data plane.

The following example shows how to display the label table and FEC information on the data plane.

```
vyatta@vyatta:~$ show dataplane mpls label-table with-prefix
Label Space: 0
in label: exp-null, fec:ipv4, outgoing label: imp-null
in label: rtr-alrt, outgoing label: imp-null
in label: exp6-null, fec:ipv6, outgoing label: imp-null
in label: 52480, fec:ipv4 192.166.4.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 52481, fec:ipv4 3.3.3.3/32
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 52482, fec:ipv4 192.166.8.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 52484, fec:ipv4 10.0.0.0/8
        nexthop via 192.168.252.253, dp0s4
```

```
in label: 52485, fec:ipv4 192.166.2.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 53120 (local), fec:ipv4 192.166.4.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 53121 (local), fec:ipv4 192.166.8.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 53128 (local), fec:ipv4 1.1.1.1/32
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: 52490
        nexthop via 192.166.1.1, dp0p1s3, outgoing label: imp-null
in label: 53129 (local), fec:ipv4 3.3.3.3/32
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
in label: 53130 (local), fec:ipv4 192.166.2.0/30
        nexthop via 192.166.3.2, dp0p1s1, outgoing label: imp-null
```

show dataplane statistics mpls

Displays the MPLS statistics on the data plane.

Syntax:

`show dataplane statistics mpls`

Operational mode

Use this command to display the MPLS statistics on the data plane.

The following example shows how to display the MPLS statistics on the data plane.

```
vyatta@vyatta:~$ show dataplane statistics mpls
MPLS Interfaces
dp0p1s11, ifindex: 10
        in bytes: 12798
        in unicast packets: 137
        in errors: 0
        label lookup failures: 0
        out bytes: 0
        out unicast packets: 0
        out errors: 0
        out fragmented packets: 0
dp0p1s12, ifindex: 11
        in bytes: 11376
        in unicast packets: 112
        in errors: 0
        label lookup failures: 0
        out bytes: 12798
        out unicast packets: 137
        out errors: 0
        out fragmented packets: 0
```

show mpls cross-connect

Displays the MPLS cross-connection information.

Syntax:

`show mpls cross-connect`

Operational mode

Use this command to display the MPLS cross-connect association of the mpls-in-segment-table and the mpls-out-segment table.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

The following example shows how to display the MPLS cross-connect association of the mpls-in-segment-table (in intf: dp0p1s1) and the mpls-out-segment table (out intf: dp0s4).

```
vyatta@vyatta:~$ show mpls cross-connect
     Cross connect ix: 3, in intf: dp0p1s1 in label: 52484 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, out intf: dp0s4, out label: N/A
    Nexthop addr: 192.168.252.253        cross connect ix: 3, op code: Pop
```

show mpls forwarding

Displays the MPLS forwarding table.

Syntax:
`show mpls forwarding`

Operational mode

Use this command to display the MPLS forwarding table that is used for imposing IPv4 traffic into an MPLS LSP.

The `show mpls forwarding-table` command displays the following information:

Output field	Description
FEC	Forwarding Equivalence Class
Nexthop	Next-hop address
Out-Label	Output label pushed onto the MPLS packet
Out-Intf	Egress interface name

The following example shows how to display the MPLS forwarding table:

```
vyatta@vyatta:~$ show mpls forwarding-table
Codes: > - selected FTN, p - stale FTN, B - BGP FTN, K - CLI FTN
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN,
       I - IGP-Shortcut, U - unknown FTN
   FEC                      Nexthop                Out-Label Out-Intf
>L 1.1.1.1/32               192.166.3.2            52490     dp0p1s1
                            192.166.1.1            3         dp0p1s3
>L 3.3.3.3/32               192.166.3.2            3         dp0p1s1
>L 192.166.2.0/30           192.166.3.2            3         dp0p1s1
>L 192.166.4.0/30           192.166.3.2            3         dp0p1s1
>L 192.166.8.0/30           192.166.3.2            3         dp0p1s1
```

show mpls in-segment-table

Displays the segment table for the incoming packet.

Syntax:
`show mpls in-segment-table`

Operational mode

Use this command to display the segment table for the incoming packet and show the incoming label-to-NHLFE mapping.

The `show mpls in-segment-table` command displays the following information:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Output field	Description
Owner:	Protocol, such as LDP
# of pops:	Number of pops
fec:	IP prefix for the FEC
ILM-ID:	Internal index of the label table
Cross connect ix:	Internal cross connect index
in intf:	Ingress router
in label:	Label on the packet from the ingress router
out-segment ix:	Internal out segment index
Owner:	Protocol, such as LDP
Persistent:	Yes or No
Admin Status:	Up or Down
Oper Status:	Up or Down
Out-segment with ix:	Internal index of the out segment
owner:	Protocol owner
out intf:	Egress router
out label:	Label on the packet for the next router (N/A for the egress router)
Nexthop addr:	Address of the next-hop router
cross connect ix:	Internal index of the cross connect
op code:	Swap or Pop

The following example shows how to display the incoming segment table and the incoming label-to-NHLFE mapping by using the `show mpls in-segment-table` command on the intermediate LSR.

```
vyatta@vyatta:~$ show mpls in-segment-table
  Owner: LDP, # of pops: 1, fec: 3.3.3.3/32, ILM-ID: 23
     Cross connect ix: 1, in intf: dp0p1s3 in label: 52481 out-segment ix: 12
      Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Up
       Out-segment with ix: 12, owner: LDP, out intf: dp0p1s1, out label: 3
    Nexthop addr: 192.166.3.2        cross connect ix: 1, op code: Swap

  Owner: LDP, # of pops: 1, fec: 10.0.0.0/8, ILM-ID: 17
     Cross connect ix: 3, in intf: dp0p1s1 in label: 52484 out-segment ix: 4
      Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Up
       Out-segment with ix: 4, owner: LDP, out intf: dp0s4, out label: N/A
    Nexthop addr: 192.168.252.253        cross connect ix: 3, op code: Pop
```

show mpls interface

Displays all MPLS-enabled interfaces.

Syntax:

`show mpls interface`

Operational mode

Use this command to display all the interfaces that are enabled for label switching and the corresponding label range that is in use for local labels.

The following example shows how to display all the MPLS- enabled interfaces.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@vyatta:~$ show mpls interface
Interface lo
  Label switching is disabled
Interface dp0s4
  Label switching is disabled
Interface dp0p1s1
  Label switching is enabled with label-space 0
   minimum label value configured is 16
   maximum label value configured is 1048575
Interface dp0p1s2
  Label switching is enabled with label-space 0
   minimum label value configured is 16
   maximum label value configured is 1048575
Interface dp0p1s3
  Label switching is enabled with label-space 0
   minimum label value configured is 16
   maximum label value configured is 1048575

Total number of mpls interface is 3
```

show mpls label-table [longer-prefixes]

Displays the MPLS label table.

Syntax:

`show mpls label-table` [*x.x.x.x/x*] [longer-prefixes]

x.x.x.x/x

The prefix or address parameter. Specify an IPv4 prefix to display any FEC that matches the supplied prefix.

Operational mode

Use this command to display the MPLS label table that is used to forward MPLS-encapsulated traffic.

The following example shows how to display the MPLS label table that is used to forward MPLS-encapsulated traffic.

```
vyatta@vyatta:~$ show mpls label-table
Codes: > - selected ILM, p - stale ILM, K - CLI ILM, T - MPLS-TP

Code FEC                 In-Lbl    Out-Lbl    In-Intf          Out-Intf         Nexthop
 >   3.3.3.3/32          52481     3          dp0p1s3          dp0p1s1          192.166.3.2
 >   10.0.0.0/8          52484     N/A        dp0p1s1          dp0s4
 192.168.252.253
 >   192.166.4.0/30      52480     3          dp0p1s3          dp0p1s1          192.166.3.2
 >   10.0.0.0/8          52484     N/A        dp0p1s3          dp0s4
 192.168.252.253
 >   192.166.8.0/30      52482     3          dp0p1s3          dp0p1s1          192.166.3.2
 >   192.166.2.0/30      52485     3          dp0p1s3          dp0p1s1          192.166.3.2
```

The following example shows how to display the MPLS label table using the longer-prefixes option.

```
vyatta@vyatta:~$ show mpls label-table 10.10.8.0/22 longer-prefixes
Codes: > - selected ILM, p - stale ILM, K - CLI ILM, T - MPLS-TP

Code FEC                 In-Lbl    Out-Lbl    In-Intf          Out-Intf         Nexthop
 >   10.10.9.0/24        53125     53123      dp0p1s11         dp0p1s12         10.10.12.7
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
>   10.10.8.0/24        53124     3         dp0p1s11         dp0p1s12         10.10.12.7
>   10.10.10.0/24       53134     3         dp0p1s12         dp0p1s11         10.10.11.5
```

show mpls out-segment-table

Displays the out segment table.

Syntax:

```
show mpls out-segment-table
```

Operational mode

Use this command to display the segment table that includes the label assignment for the MPLS traffic leaving the LSR.

The `show mpls out-segment-table` command displays the following information:

Output field	Description
Out-segment with ix:	Internal index of the out segment
owner:	Protocol, such as LDP
out intf:	Name of the interface from which traffic is leaving
out label:	Label on the packet for the next router, or N/A if this is the egress router.
Nexthop addr:	Address of the next-hop router
cross connect ix:	Internal index of the cross connect
op code:	Push

The following example shows how to display the label assignment for the MPLS traffic that is leaving the LSR.

```
vyatta@vyatta:~$ show mpls out-segment-table
       Out-segment with ix: 12, owner: LDP, out intf: dp0p1s1, out label: 3
    Nexthop addr: 192.166.3.2        cross connect ix: 1, op code: Push

       Out-segment with ix: 4, owner: LDP, out intf: dp0s4, out label: N/A
    Nexthop addr: 192.168.252.253        cross connect ix: 3, op code: Push
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Label Distribution Protocol

LDP overview

When used to create LSP tunnels, Label Distribution Protocol (LDP) allows a set of destination IP prefixes (known as a Forwarding Equivalence Class or FEC) to be associated with an LSP.

Each Link State Router (LSR) establishes a peer relationship with the neighboring LDP-enabled routers and exchanges label mapping information. This label mapping information is stored in an LDP database on each LSR. When an LSR determines that one of the peers is the next-hop for a FEC, the LSR uses the label mapping information from the peer to set up an LSP that is associated with the FEC.

The devices advertise their loopback addresses to their LDP peers as a 32-bit prefix-type FEC. When an LSR installs a label for a FEC, it also creates an MPLS tunnel route, which is then made available to routing applications. This allows each router to potentially be an ingress LER for an LSP whose destination is the device's loopback address.

The result of an LDP configuration is a full mesh of LSPs in an MPLS network, with each LDP-enabled router a potential ingress, transit, or egress LSR, depending on the destination.

The system supports LDP for the configuration of non-traffic-engineered tunnel LSPs in an MPLS network. LDP is described in *RFC 5036*.

The AT&T Vyatta vRouter implementation supports the following aspects of LDP:

- Liberal label retention—Each LSR sends its peers Label Mapping messages, which map a label to a FEC. The peer LSR receiving these messages retain all of the mappings, even though they may not actually be used for data forwarding.
- Unsolicited label advertisement—The LSR sends Label Mapping messages to its LDP peers even though they did not explicitly request them.
- Ordered label distribution—The LSR sends a Label Mapping message to its peers only when it knows the next hop for a FEC, or is itself an egress LER for the FEC. When an LSR does not know the next hop for a FEC, and is not an egress LER for the FEC, it waits until a downstream LSR sends it a Label Mapping message for the FEC. At this point, the LSR can send Label Mapping messages for the FEC to its peers. This allows label mappings to be distributed, in an orderly fashion, starting from the egress LER and progressing upstream.

The LDP label space ID has a default value of zero which improves interoperability with routers from other vendors.

LDP terminology

Before implementing LDP, familiarize yourself with the following key terms and definitions.

Apply Current Route	Compare the current route with the received label mappings and install any downstream mappings, as appropriate.
Current Route	Current next hop info for a FEC. The current route may not be applied immediately to the current label mapping as a result of LWD at ingress.
Downstream mapping (DM)	Represents the label mapping received from a downstream peer for a FEC.
FEC	Forwarding Equivalency Class. Each FEC is a destination IP address for an LDP tunnel.
LDP	Label Distribution Protocol.
Label mapping	LDP message that indicates the label to be used for an FEC from the peer.
LSP	Label Switched Path.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

LWD	Label Withdrawal Delay.
Route event	Update from the routing table to LDP.
Upstream Mapping	Represents the label mapping sent to an upstream peer for a FEC.

Configuring LDP on an interface

To establish LDP sessions and exchange labels with a peer, LDP must be enabled on the neighbors' interfaces.

To use LDP, configure a loopback address with a 32-bit mask on the LSR. The first loopback address configured on the device is used in its LDP identifier. When the loopback address used in the LDP identifier is removed, all LDP functions on the LSR are shut down. LDP sessions between the LSR and its peers are terminated, and LDP-created tunnels are removed. When other loopback interfaces are configured on the device, the lowest-numbered loopback address is used as a new LDP identifier. LDP sessions and tunnels are set up using this new LDP identifier.

Configure LDP on the same set of interfaces that IGP routing protocols such as OSPF are enabled.

To configure LDP on an interface, perform the following steps.

1. Enable MPLS LDP on an interface in configuration mode.

   Example:

   ```
   vyatta@R1# set protocol mpls-ldp discovery interfaces interface dp0p1s1 address-family
    ipv4
   ```

2. Verify the configuration of the interface in operational mode.

   Example:

   ```
   vyatta@vyatta:~$ show mpls ldp interface dp0p1s1
   Interface        LDP Identifier           LDP Enabled Version Merge Capability
   dp0p1s1          2.2.2.2:0                Enabled     IPv4    Merge capable
   ```

The following example displays MPLS LDP configured on three interfaces.

```
vyatta@vyatta:~$ show mpls ldp interface
Interface        LDP Identifier           LDP Enabled Version Merge Capability
lo               2.2.2.2:0                Disabled            N/A
dp0s4            2.2.2.2:0                Disabled            N/A
dp0p1s1          2.2.2.2:0                Enabled     IPv4    Merge capable
dp0p1s2          2.2.2.2:0                Enabled     IPv4    Merge capable
dp0p1s3          2.2.2.2:0                Enabled     IPv4    Merge capable
```

LDP outbound FEC filtering

LDP outbound FEC filtering allows LDP to perform outbound filtering for label advertisement. It gives you the ability to control which FECs can be advertised and to which LDP neighbors. It also reduces the number of labels distributed to neighbors and the number of messages exchanged with peers. Through this feature, LDP scalability and convergence, security, and performance are improved.

LDP performs a hop-by-hop or dynamic path setup in an MPLS network by assigning and distributing labels to routes learned from the underlying IGP routing protocols. By default, LDP distributes all FECs that are learned locally or from LDP neighbors to all other LDP neighbors. When this behavior is not desired, you can configure LDP to perform outbound FEC filtering.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Outbound filtering is achieved by creating a prefix list that specify prefixes whose label mappings can be distributed. The prefix list is applied globally to all the LDP neighbors. The FECs permitted by the prefix list are accordingly distributed to the specified LDP neighbor or to all LDP neighbors.

Prerequisites

MPLS and LDP protocols must be enabled on the router to use this feature.

Configuring LDP outbound FEC filtering

Configures a prefix list.

MPLS LDP must be enabled.

Configure a prefix list to avoid advertising the 10.0.1.0/24 prefix.

1. Edit the prefix list.

   Example:

   ```
   vyatta@R1# edit policy route prefix-list label-policy10
   vyatta@R1# set rule 10 action deny
   vyatta@R1# set rule 10 prefix 10.0.1.0/24
   vyatta@R1# set rule 20 action permit
   vyatta@R1# set rule 20 prefix 0.0.0.0/0
   vyatta@R1# set rule 20 le 32
   ```

2. Configure LDP outbound filtering

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp address-family ipv4 label-policy advertise prefix-
   list label-policy10
   ```

LDP Hello interval and Hello hold timeout timers

The LDP Hello interval and Hello hold timeout timers are used to establish Hello adjacency between peers. The Hello interval is the time period between which the LSR sends out Hello messages and the Hello hold timeout is the amount of time that the sending LSR maintains its record of Hellos from the receiving LSR without receipt of another Hello message.

The Hello interval and Hello hold timeout timer values can be obtained from the global default values or configured globally on a router. The Hello hold timeout timer value can also be configured through an interface. When configuring these values the following constraints must be followed:

- The Hello interval value must be less than 1200
- The Hello hold timeout value must be less than 3600
- The Hello hold timeout value must be greater than or equal to 3 times the Hello interval value

The values can be set that determine the values used on the configured router and values sent to adjacent peers for their configuration. The following sections describe how to set the values:

- Setting the LDP Hello interval values
- Setting the LDP Hello hold time sent to adjacent LSRs
- Determining the LDP Hello hold time on an MPLS interface

Precedence of the Hello hold time settings

The precedence of the Hello hold time value settings is as follows:

Determining the LDP hold time on an MPLS interface:

An MPLS interface uses the LDP Hello hold time to determine how long it waits for its LDP peers to send a Hello message.

# For link LDP sessions - In this case, the wait time is determined by any one of the below criteria.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

1. When the Hello hold time is set per-interface, that value is used.
2. When the Hello hold time is not set per-interface, the hold time in the received message is used.
3. When the Hello hold time in the received message is zero (0), the default value of 15 seconds is used.

Precedence of the Hello interval settings

The precedence of the Hello interval value settings is as follows:

- For link LDP sessions - the LDP Hello interval can be set globally which applies to all LDP interfaces or on a per-interface basis. The LDP Hello interval values in LDP link sessions are determined by the following procedure in the order described below.

1. When the Hello interval is set per-interface, that value is used.
2. When the Hello interval is not set per-interface, then the value set for LDPs globally is used.
3. When the Hello interval is not set either globally or per-interface, the global default value is used. When a Hello adjacency already exists, the adjacency remains up and any new configured interval takes effect upon the expiration of the current Hello interval timer. Consequently, the next and subsequent Hello messages are sent at the new interval.

Setting LDP Hello interval values

Set the LDP Hello interval value to configure the time period between which the LSR sends out Hello messages

1. To change the global link and target intervals for LDP Hello messages, and configure a link interval for an interface, enter the following command.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s1 hello-
   interval
   ```

> In the following example, the hello interval set to 10. The default value is 5. You can enter an integer from 5 through 1200.
>
> ```
> vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s1 hello-interval 10
> ```

Setting LDP Hello hold timeout

To change the default hold time included in LDP Hello messages, perform the following step.

1. Change the hold time for link sessions on the interface. In this example a hold time of 30 is specified. You can enter an integer between 15 and 3600.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s1 hello-holdtime
    30
   ```

Setting hold time and interval globally

Set LDP hold time and Hello interval globally.

For an LDP session between routers, you must configure LDP on an interface to allow the device to advertise its loopback interface to the peers.

To set the hold time and interval included in LDP Hello messages globally perform the following steps.

1. Set the hold time globally.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp discovery interfaces hello-holdtime 30
   ```

2. Set the interval globally.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp discovery interfaces hello-interval 10
   ```

Configuring LDP message authentication

Configuration of an authentication key on a per LDP session basis is supported to protect against spoofed TCP segments in a connection stream.

The LDP session can be to an adjacent peer (basic discovery) . You must configure both sides of an LDP peer link.

LDP authentication is based upon the TCP MD5 signature option specified in *RFC 2385*. This RFC defines a new TCP option for carrying an MD5 digest in a TCP segment.

1. To configure LDP message authentication, enter the following command.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp neighbors neighbor <ip address> md5-password <text>
   ```

The following example shows how to configure an LDP message authentication key named early for the neighbor at IP address 10.10.10.3.

```
vyatta@R1# set protocols mpls-ldp neighbors neighbor 10.10.10.3 md5-password early
```

Resetting LDP neighbors

You can terminate and re-establish an MPLS LDP neighbor session when at least one LDP Hello adjacency exists with the peer. When the LDP session terminates, the database associated with the LDP session is also cleared.

The LDP sessions are automatically re-established when at least one Hello adjacency exists with the neighbor, and LDP configuration remains unchanged.

When an LDP session is terminated as a result of the `reset mpls ldp` command, the AT&T device generates a notification message for the neighbor. The other end of the LDP session detects this reset operation and tries to re-establish the session.

Note: Resetting an LDP session that is not in an operational state has no impact.

1. Enter the `reset mpls ldp session all` command to terminate all LDP sessions.

   Example:

   ```
   vyatta@vyatta:~$ reset mpls ldp session all
   ```

   Info: In this example, both the link and targeted LDP sessions are terminated. When the **`all`** option is specified instead of a peer address, all LDP sessions on the AT&T device are reset.

2. Enter the `reset mpls ldp session <ip address>` command to terminate the LDP sessions with the neighbor.

   Example:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@vyatta:~$ reset mpls ldp session 10.234.123.64
```

Info: In this example, the link LDP sessions with neighbor 10.234.123.64 are terminated.

When the session re-establishes, the session-specific information is re-learned from its peer:

- LDP downstream and upstream label database displayed by the `show mpls ldp downstream` and `show mpls ldp upstream` command
- LDP label switched path displayed by the `show mpls ldp lsp` command
- LDP peer displayed by the `show mpls ldp neighbor` command
- LDP FECs learned from the resetting neighbor sessions displayed by the `show mpls ldp fec` command FECs are not cleared immediately but are marked that no LDP session exists.

Validating LDP session reset

You can check the LDP session specific parameters to validate that a session has been successfully reset.

- The LDP session state transitions from `OPERATIONAL` to `NON_EXISTANT` upon clearing it. However, the states may quickly transition. In this case, the `show mpls ldp session` command shows the current state.
- The LDP session-specific database is cleaned upon resetting the LDP session.
- The TCP port number on the active end of the LDP session may have been changed once the LDP session comes up after the reset; the TCP port number before and after the reset may be different.
- Syslog logs the event of a LDP session going down and then coming back up as a result of resetting the LDP session. Use the command `show log` to view the syslog events.

The following example shows the use of the `show mpls ldp session` command to validate the LDP session reset.

```
vyatta@vyatta:~$ show mpls ldp session
Peer IP Address           IF Name         My Role State         KeepAlive
5.5.5.5                   dp0p1s3         Passive OPERATIONAL    30
```

MPLS LDP-IGP synchronization

MPLS LDP-IGP synchronization provides a means to synchronize LDP and IGPs to minimize MPLS packet loss.

MPLS LDP-IGP synchronization also provides the following benefits:

- Provides a means to disable LDP-IGP synchronization on interfaces that you do not want enabled
- Allows you to globally enable LDP-IGP synchronization on each interface associated with an IGP Open Shortest Path First (OSPF)

MPLS LDP-IGP synchronization may be enabled on an interface. LDP determines convergence (receipt of all labels) for a link through one of two methods:

- Receive Label silence mechanism
- End Of Lib mechanism (*RFC 5919*)

The following figure provides an example of LDP-IGP synchronization.

Figure 5: Example with LDP-IGP synchronization

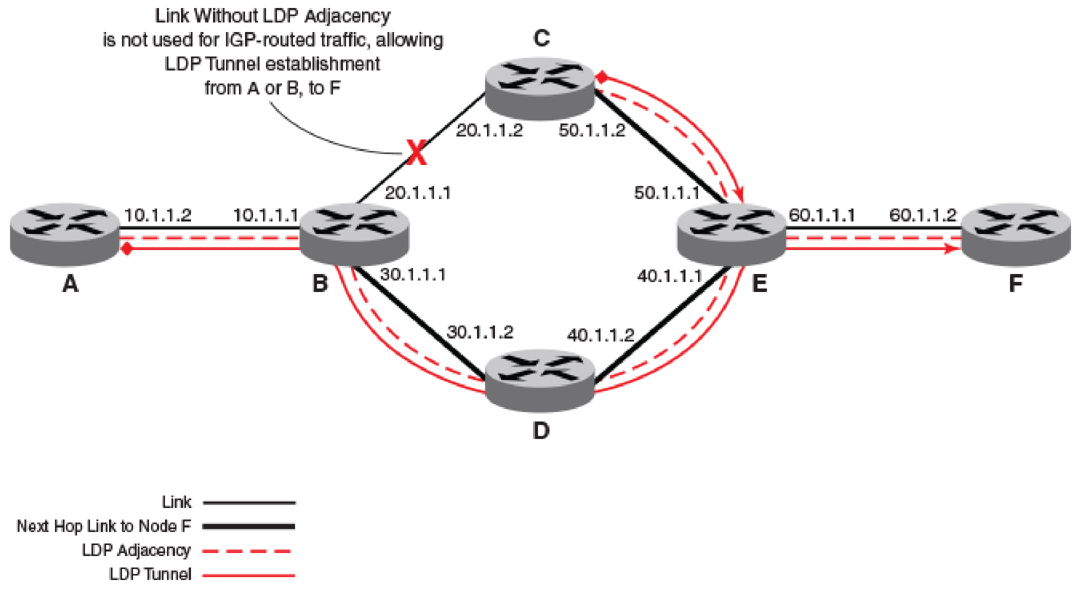

When the LDP peer is not reachable, the IGP establishes the adjacency to enable the LDP session to be established. When an IGP adjacency is established on a link but LDP-IGP Synchronization is not yet achieved or is lost, the IGP advertises the maximum metric (max-metric) on the link.

Configuration considerations

- Affects IPv4 metrics only

### Configuring MPLS LDP-IGP synchronization

Configuring MPLS LDP-IGP includes the following tasks:

# Configuring MPLS LDP-IGP synchronization with OSPF interfaces (required)

# Selectively disabling MPLS LDP-IGP synchronization from some OSPF interfaces (optional)

# Verifying MPLS LDP-IGP synchronization with OSPF (optional)

Enabling MPLS LDP-IGP synchronization on an interface

You can enable LDP-IGP synchronization on an interface that belongs to an OSPF process and override global LDP-IGP synchronization.

Perform the following steps.

1. Enable LDP-IGP synchronization with OSPF on the interface.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s1 igp-
   synchronization-delay <delay>
   ```

2. Verify LDP-IGP synchronization with OSPF on the interface.

   Example:

   ```
   vyatta@vyatta:~$ show ip ospf interface dp0p1s1
   dp0p1s1 is up, line protocol is up
     Internet Address 192.166.2.1/30, Area 0.0.0.0, MTU 1500
     Process ID 0, routing-instance (default), Router ID 1.1.1.1, Network Type BROADCAST,
    Cost: 10
     Transmit Delay is 1 sec,  State DR, Priority 1, TE Metric 10
   ```

```
  LDP-OSPF Sync configured
    Holddown timer not configured
  Designated Router (ID) 1.1.1.1, Interface Address 192.166.2.1
  Backup Designated Router (ID) 3.3.3.3, Interface Address 192.166.2.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Neighbor Count is 1, Adjacent neighbor count is 1
  Hello received 197 sent 205, DD received 3 sent 4
  LS-Req received 1 sent 1, LS-Upd received 71 sent 32
  LS-Ack received 9 sent 37, Discarded 0
  No authentication
```

3. Verify LDP-IGP synchronization with LDP.

   Example:

```
vyatta@vyatta:~$ show mpls ldp igp sync
lo is up, line protocol is up
  LDP not configured; LDP-IGP Synchronization not enabled.
dp0s4 is up, line protocol is up
  LDP not configured; LDP-IGP Synchronization not enabled.
dp0p1s1 is up, line protocol is up
  LDP configured; LDP-IGP Synchronization enabled.
  Session IP Address : 3.3.3.3
    Sync status: Achieved
    Delay timer: Configured, 15 seconds, Not Running
dp0p1s2 is up, line protocol is up
  LDP not configured; LDP-IGP Synchronization not enabled.
dp0p1s3 is up, line protocol is up
  LDP configured; LDP-IGP Synchronization enabled.
  Session IP Address : 2.2.2.2
    Sync status: Achieved
    Delay timer: Configured, 15 seconds, Not Running
dp0p1s4 is up, line protocol is up
  LDP not configured; LDP-IGP Synchronization not enabled.
lo1 is up, line protocol is up
  LDP not configured; LDP-IGP Synchronization not enabled.
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Configuring LDP

The minimum configuration required to enable MPLS with LDP is to enable LDP on an interface.

Configure the interfaces on which LDP should be activated.

To create a configuration that supports LDP, perform the following steps in configuration mode.

1. Enable MPLS with LDP on a specific interface.

   Example:

```
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s1 address-family
 ipv4
vyatta@R1# commit
```

2. Repeat Step 1 for all interfaces on which LDP is be enabled.
3. Display the MPLS LDP status for a specific interface.

   Result:

```
vyatta@R1# exit
vyatta@R1:~$ show mpls ldp discovery dp0p1s1
Status                  : Enabled
Version                 : IPv4
Primary IP Address      : 192.166.3.1
Interface Type          : Ethernet
Label Merge Capability  : Merge Capable
Hold Time               : 15
Hello Interval          : 5
Targeted Hello Interval : 15
Targeted Hold Time      : 45
Keepalive Interval      : 10
Keepalive Timeout       : 30
Advertisement Mode      : Downstream Unsolicited
Label Retention Mode    : Liberal
Multicast Hellos        : Enabled
Max PDU Length          : 4096
```

4. Display all interfaces configured for LDP discovery.

   Example:

```
vyatta@R1:~$ show mpls ldp discovery
Interface        LDP Identifier           LDP Enabled Version Merge Capability
lo               2.2.2.2:0                Disabled            N/A
dp0s4            2.2.2.2:0                Disabled            N/A
dp0p1s1          2.2.2.2:0                Enabled     IPv4    Merge capable
dp0p1s2          2.2.2.2:0                Enabled     IPv4    Merge capable
dp0p1s3          2.2.2.2:0                Enabled     IPv4    Merge capable
```

In the following example, three interfaces are enabled for MPLS LDP and the `show mpls ldp discovery` command is used to display all interfaces configured for LDP discovery.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s1 address-family ipv4
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s2 address-family ipv4
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s3 address-family ipv4
vyatta@R1# commit
vyatta@R1# exit
vyatta@R1:~$ show mpls ldp discovery
Interface        LDP Identifier           LDP Enabled Version Merge Capability
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
lo                2.2.2.2:0                Disabled           N/A
dp0s4             2.2.2.2:0                Disabled           N/A
dp0p1s1           2.2.2.2:0                Enabled     IPv4   Merge capable
dp0p1s2           2.2.2.2:0                Enabled     IPv4   Merge capable
dp0p1s3           2.2.2.2:0                Enabled     IPv4   Merge capable
```

LDP Graceful Restart

LDP Graceful Restart (GR) helps minimize MPLS traffic loss when an LDP component is restarting in a router that is capable of preserving its MPLS forwarding states across restart. LDP GR is based on RFC 3478 (Graceful Restart mechanism for Label Distribution Protocol).

LDP GR works between a router and its neighbor and its capability must be advertised when sending an LDP Initialization message.

The router can also support LDP GR in helper-only mode. In this mode, a router does not preserve its forwarding entries on a LDP GR restart. It indicates to its peers that forwarding state is not preserved by sending an initialization message with the Reconnect Time and the Recovery Time set to zero (0) in FT session TLV. However, it can help a neighboring router recover its forwarding entries when the neighbor is going through restart.

An LDP GR enabled router goes into helper-only mode (GR helper) when any of the following events occur on the neighboring routers.

- Remove and re-add of the MPLS configuration
- TCP communication broken (such as, session keepalive timer expires)
- UDP communication broken (for example, an adjacency goes down)
- Restarting LDP component by disabling and enabling the loopback
- Restarting a LDP session by issuing the `reset mpls ldp session` command

In helper-only mode, the LDP GR procedure works at the session level. Any of the previous events causes the helper to detect session down and start the GR procedure. The operation of the GR helper is the same independent of what has happened on the restarting LSR that triggers the GR procedure.

When LDP GR is enabled on a router, the configuration does not apply to the current sessions. The LDP GR configuration is applied for the new sessions brought up after the configuration is added.

Configuring LDP GR helper

By default LDP GR is disabled. You can globally enable it. When LDP GR enabled, the router waits until it receives an LDP Initialization message from its neighbor to know whether it must delete its states or start the LDP GR recovery procedure. Also, LDP GR is applicable to all LDP sessions regardless of the adjacency type that exists between the neighbors.

Note: The following command only takes effect on newly created sessions. For existing sessions, it is required that the sessions be restarted for the new configuration to take effect.

1. Configure MPLS LDP GR helper mode.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp graceful-restart helper-enable
   ```

LDP session Keepalive timeout configurations

Keepalive messages are used to help maintain the integrity of the LDP session.

After an LDP session is established, an LSR maintains the integrity of the session by sending Keepalive messages. The Keepalive timer for each peer session resets whenever it receives any LDP protocol message or a

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Keepalive message on that session. When the Keepalive timer expires, LDP concludes that the TCP connection is bad or the peer is dead and terminates the session.

Setting the Keepalive timeout

Use session-ka-timeout to set the Keepalive interval or the time interval at which the session Keepalive message is sent when no other LDP protocol message is sent to the LDP peer.

To configure the Keepalive timeout, perform the following step.

1. Set the Keepalive timeout in the range of 30 through 3600 seconds.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp neighbors session-ka-holdtime 60
   ```

Setting the Keepalive intervals

Use session-ka-intervals to configure the number of intervals after which the session is terminated when no session Keepalive or other LDP protocol message is received from the LDP peer.

To configure the Keepalive interval, perform the following step.

1. Set the Keepalive interval in the range of 10 through 1200 seconds.

   Example:

   ```
   vyatta@R1#  set protocols mpls-ldp neighbors session-ka-interval 20
   ```

Configurable LDP LSR ID

LDP uses LDP messages to communicate between LDP peers for the correct functioning of LDP. All LDP messages contains a LDP header which is composed of LDP version, length of message, LDP ID, and is followed by a message. The LDP ID for LDP is composed of the LSR-ID and label space. The LSR ID is the first available loopback interface address. However, you can specify an IP address of your choice to use as the LSR ID for the LDP identifier.

By default, AT&T routers select the first valid and operationally UP IP address among all the enabled loopback interfaces as LSR ID for LDP. When the IP address or loopback interface that is used as the LSR ID goes down, LDP selects the next operationally UP IP address among all enabled loopback interface as the LSR ID. Otherwise, LDP will be down.

When no valid IP address is available to be selected as the LSR ID, LDP continues to remain disabled until a valid IP address is configured on an enabled loopback interface.

Figure 6: LDP Header format

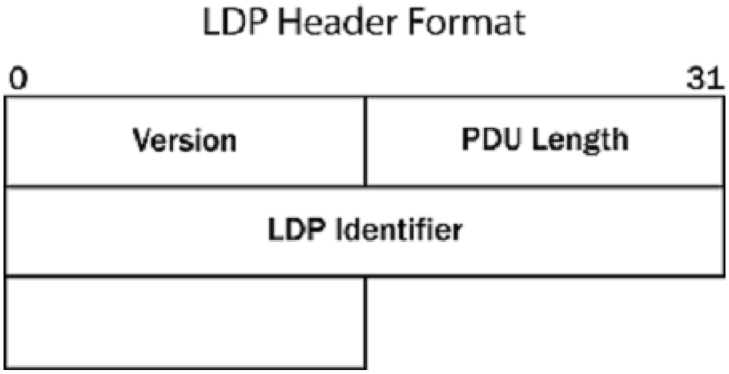

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

LDP uses a configured IP address as the LSR ID only when this IP address is configured on one of the enabled loopback interfaces. After you configure an LSR ID with a valid IP address, LDP must use the configured value as the LSR ID and restarts to use the new address.

When this IP address is not configured in the enabled state on any of the loopback interfaces, LDP continues in the disabled state. LDP is enabled as soon as this IP address is configured on one of the enabled loopback interfaces.

When you disable the feature, the LSR ID selection procedure falls back to default behavior of selecting an LSR ID for LDP when LDP is enabled.

Limitations

- You cannot configure value 0.0.0.0. If you try to configure the feature with this value, the feature rejects the configuration.
- You can only configure IPv4 addresses for the LSR ID.

Configuring the LDP LSR ID

By default, the LSR-ID is the first available loopback interface address. However, you can specify an IP address of your choice to use as the LSR-ID for the LDP identifier.

Ensure that the LSR ID IP address is an operationally UP IP address on an enabled loopback interface.

1. Enter the following command to configure an LSR ID.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp lsr-id <ip address>
   ```

   Info: The configured IP address selected as the LSR ID for LDP is an operationally UP IP address on an enabled loopback interface. You can configure only an IPv4 address.

The following example shows how to configure the LDP router ID as 1.1.1.1.

```
vyatta@R1# set protocols mpls-ldp lsr-id 1.1.1.1
```

Sample LDP Configurations

This section presents examples of typical MPLS LDP configurations.

Configuring MPLS LDP example

How to create a basic MPLS LDP configuration

MPLS LDP must be enabled on interfaces before MPLS LDP is supported.

The minimum configuration required to enable MPLS with LDP is to configure the interfaces on which LDP should be activated:

1. Enable LDP on the interface.

   Example:

   ```
   vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s9 address-family
    ipv4
   vyatta@R1# commit
   vyatta@R1# exit
   ```

2. Verify the MPLS LDP configuration.

   Example:

   ```
   vyatta@R1:~$ show mpls ldp session
   Peer IP Address        IF Name        My Role    State        KeepAlive
   4.4.4.4                dp0p1s9        Passive    OPERATIONAL  30
   ```

Minimum MPLS LDP example

Minimum configuration example for MPLS LDP

MPLS LDP must be enabled on interfaces before MPLS LDP is supported.

This configuration example is for MPLS LDP on LSR1, LSR2, and LSR3 in the following topology:

LSR1 ------ LSR2 ------- LSR3

1. Enter the following configuration commands on LSR1.

   Example:

   ```

    set interfaces dataplane dp0p1s1 address 10.1.2.1/30
    set interfaces loopback lo1 address 1.1.1.1/32
    set protocols mpls-ldp address-family ipv4 transport-address 1.1.1.1
    set protocols mpls-ldp discovery interfaces interface dp0p1s1 address-family ipv4
    set protocols mpls-ldp lsr-id 1.1.1.1
    set protocols ospf area 0.0.0.0 network 1.1.1.1/32
    set protocols ospf area 0.0.0.0 network 10.1.2.0/30
    set protocols ospf parameters router-id 1.1.1.1
    commit
   ```

2. Enter the following configuration commands on LSR2.

   Example:

   ```

    set interfaces dataplane dp0p1s1 address 10.1.2.2/30
    set interfaces dataplane dp0p1s2 address 10.1.3.1/30
    set interfaces loopback lo1 address 2.2.2.2/32
    set protocols mpls-ldp discovery interfaces interface dp0p1s1 address-family ipv4
   ```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
set protocols mpls-ldp discovery interfaces interface dp0p1s2 address-family ipv4
set protocols mpls-ldp lsr-id 2.2.2.2
set protocols ospf area 0.0.0.1 network 2.2.2.2/32
set protocols ospf area 0.0.0.0 network 10.1.2.0/30
set protocols ospf area 0.0.0.0 network 10.1.3.0/30
set protocols ospf parameters router-id 2.2.2.2
commit
```

3. Enter the following configuration commands on LSR3.

   Example:

```
set interfaces dataplane dp0p1s1 address 10.1.3.2/30
set interfaces loopback lo1 address 3.3.3.3/32
set protocols mpls-ldp discovery interfaces interface dp0p1s1 address-family ipv4
set protocols mpls-ldp lsr-id 3.3.3.3
set protocols ospf area 0.0.0.0 network 3.3.3.3/32
set protocols ospf area 0.0.0.0 network 10.1.3.0/30
set protocols ospf parameters router-id 3.3.3.3
commit
```

Verify the MPLS LDP configuration by entering the following command on LSR2.

```
run show mpls ldp session
Peer IP Address            IF Name          My Role State          KeepAlive
1.1.1.1                    dp0p1s1          Passive OPERATIONAL    30
3.3.3.3                    dp0p1s2          Passive OPERATIONAL    30
```

LDP Commands

List of the supported MPLS LDP operational commands and configuration commands

The supported MPLS LDP commands are grouped in the following sections:

- MPLS LDP operational commands
- MPLS LDP configuration commands

LDP operational commands

The following clear, reset, and show commands are available with MPLS LDP:

1. clear mpls ldp statistics
2. reset mpls ldp [adjacency all | X.X.X.X] | [session all | X.X.X.X]
3. show mpls ldp
4. show mpls ldp adjacency
5. show mpls ldp advertise-labels
6. show mpls ldp discovery [<interface>]
7. show mpls ldp downstream
8. show mpls ldp fec
9. show mpls ldp graceful-restart
10. show mpls ldp igp sync
11. show mpls ldp interface [<interface>]
12. show mpls ldp isp [detail | host | prefix]
13. show mpls ldp neighbors
14. show mpls ldp routes
15. show mpls ldp session [X.X.X.X]
16. show mpls ldp statistics
17. show mpls ldp upstream

LDP configuration commands

The following global configuration commands are available with MPLS LDP:

1. protocols mpls-ldp lsr-id <ipaddr>
2. protocols mpls-ldp neighbors neighbor <ip address> md5-password <text>
3. protocols mpls-ldp neighbors session-downstream-on-demand
4. protocols mpls-ldp neighbors session-ka-holdtime 30..3600
5. protocols mpls-ldp neighbors session-ka-interval 10..1200
6. protocols mpls-ldp address-family ipv4 label-policy advertise explicit-null
7. protocols mpls-ldp address-family ipv4 label-policy advertise prefix-list <name>
8. protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode independent
9. protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode ordered
10. protocols mpls-ldp address-family ipv4 transport-address <ip address>
11. protocols mpls-ldp graceful-restart helper-enable
12. protocols mpls-ldp graceful-restart reconnect-time 10..1800
13. protocols mpls-ldp graceful-restart recovery-time 30..3600

The following interface configuration commands are available with MPLS LDP:

1. protocols mpls-ldp discovery interfaces interface <name> address-family ipv4
2. protocols mpls-ldp discovery interfaces hello-holdtime 15..3600
3. protocols mpls-ldp discovery interfaces hello-interval 5..1200
4. protocols mpls-ldp discovery interfaces interface <name> igp-synchronization-delay 3..60

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

5. protocols mpls-ldp discovery interfaces interface <name> hello-holdtime 15..3600
6. protocols mpls-ldp discovery interfaces interface <name> hello-interval 5..1200
7. protocols mpls-ldp discovery interfaces interface <name> session-ka-holdtime 30..3600
8. protocols mpls-ldp discovery interfaces interface <name> session-ka-interval 10..1200

The following logging configuration commands are available with MPLS LDP:

1. protocols mpls-ldp log dsm
2. protocols mpls-ldp log events
3. protocols mpls-ldp log fsm
4. protocols mpls-ldp log nsm
5. protocols mpls-ldp log rib
6. protocols mpls-ldp log usm
7. protocols mpls-ldp log packet address | all | hello | init | keepalive | label | notification

Logging is enabled for LDP with the following command:

monitor protocols mpls ldp enable | disable [dsm | events | fsm | nsm | packet | rib | usm]

clear mpls ldp statistics

Clears the packet statistics that are displayed by the `show mpls ldp statistics` command.

Syntax:

```
clear mpls ldp statistics
```

Operational mode

Use this command to clear the packet statistics that are displayed by the `show mpls ldp statistics` command.

monitor protocol mpls ldp enable|disable [dsm | events | fsm | nsm | packet | rib | usm]

Enables or disables the generation of debug messages that are related to MPLS LDP logs.

Syntax:

```
monitor protocol mpls ldp enable|disable dsm | events | fsm | nsm | packet | rib | usm
```

Operational mode

Use this command to enable or disable the generation of debug messages that are related to MPLS LDP logs.

The following example shows how to enable MPLS LDP Finite State Machine monitoring.

```
vyatta@vyatta:~$ monitor protocol mpls ldp enable fsm
```

As a result, the LDP debugging status is on when displayed with the `show monitoring` command.

```
vyatta@vyatta:~$ show monitoring
------------------------------
  Protocol monitoring status
------------------------------

...

LDP debugging status:
  ...
  LDP finite state machine debugging is on
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
...
```

protocols mpls-ldp address-family ipv4 label-policy advertise explicit-null

Enables an egress router to advertise an explicit null label in place of an implicit null label to the penultimate hop router.

Syntax:
`set protocols mpls-ldp address-family ipv4 label-policy advertise explicit-null`

Syntax:
`delete protocols mpls-ldp address-family ipv4 label-policy advertise explicit-null`

Syntax:
`show protocols mpls-ldp address-family ipv4 label-policy advertise explicit-null`

Configuration mode

```
protocols {
    mpls-ldp {
        address-family ipv4  {
            label-policy {
                advertise explicit-null
            }
        }
    }
}
```

Use this command to enable an egress router to advertise an explicit null label (value 0) in place of an implicit null label (value 3) to the penultimate hop router.

Use the `set` form of this command to enable an egress router to advertise an explicit null label in place of an implicit null label to the penultimate hop router.

Use the `delete` form of this command to delete the configuration that enables an egress router to advertise an explicit null label in place of an implicit null label to the penultimate hop router.

Use the `show` form of this command to display the configuration that enables an egress router to advertise an explicit null label in place of an implicit null label to the penultimate hop router.

protocols mpls-ldp address-family ipv4 label-policy advertise prefix-list

Applies the prefix list name to filter outgoing label advertisements.

Syntax:
`set protocols mpls-ldp address-family ipv4 label-policy advertise` prefix-list *name*

Syntax:
`delete protocols mpls-ldp address-family ipv4 label-policy advertise` prefix-list *name*

Syntax:
`show protocols mpls-ldp address-family ipv4 label-policy advertise` prefix-list *name*

name
The prefix list name.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Configuration mode

```
protocols {
    mpls-ldp {
        address-family {
            ipv4  {
                label-policy {
                    advertise {
                        prefix-list <name>
                    }
                }
            }
        }
    }
}
```

Use this command to apply the prefix list name to filter outgoing label advertisements.

Use the `set` form of this command to apply the prefix list name to filter outgoing label advertisements.

Use the `delete` form of this command to delete the configuration that applies the prefix list name to filter outgoing label advertisements.

Use the `show` form of this command to display the configuration that applies the prefix list name to filter outgoing label advertisements.

Example: Example

The following example shows how to apply the prefix list named list1 to filter outgoing label advertisements.

```
vyatta@R1# set protocols mpls-ldp address-family ipv4 label-policy advertise prefix-list
 list1
```

protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode independent

Enables independent label mode.

Syntax:
`set protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode independent`

Syntax:
`delete protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode independent`

Syntax:
`show protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode independent`

Configuration mode

```
protocols {
    mpls-ldp {
        address-family {
            ipv4  {
                label-policy {
                    distribution-control-mode {
                        independent
                    }
                }
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
            }
        }
    }
}
```

Use this command to enable independent label mode (the default behavior) where a label mapping to a FEC is advertised as long as a RIB entry exists for the given FEC.

Use the `set` form of this command to enable independent label mode.

Use the `delete` form of this command to delete the configuration that enables independent label mode.

Use the `show` form of this command to display the configuration that enables independent label mode.

protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode ordered

Enables ordered label mode.

Syntax:
`set protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode ordered`

Syntax:
`delete protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode ordered`

Syntax:
`show protocols mpls-ldp address-family ipv4 label-policy distribution-control-mode ordered`

Configuration mode

```
protocols {
    mpls-ldp {
        address-family {
            ipv4  {
                label-policy {
                    distribution-control-mode {
                        ordered
                    }
                }
            }
        }
    }
}
```

Use this command to enable ordered label mode where a label mapping to a FEC is advertised if this is an egress LSR or if a label binding for the FEC has been received.

Use the `set` form of this command to enable ordered label mode.

Use the `delete` form of this command to delete the configuration that enables ordered label mode.

Use the `show` form of this command to display the configuration that enables ordered label mode.

protocols mpls-ldp address-family ipv4 transport-address

Configures the ipv4 transport address used for the MPLS LDP TCP session.

Syntax:
`set protocols mpls-ldp address-family ipv4` transport-address *ip address*

Syntax:
`delete protocols mpls-ldp address-family ipv4` transport-address *ip address*

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`show protocols mpls-ldp address-family ipv4` transport-address *ip address*

ip address
The IP address.

Configuration mode

```
protocols {
    mpls-ldp {
        address-family {
            ipv4 {
                transport-address <ip address>
            }
        }
    }
}
```

Use this command to configure the ipv4 transport address used for the LDP TCP session. The address can be bound to a loopback interface or a physical interface.

Use the `set` form of this command to configure the ipv4 transport address used for the LDP TCP session.

Use the `delete` form of this command to delete the configuration of the ipv4 transport address used for the LDP TCP session.

Use the `show` form of this command to display the configuration of the ipv4 transport address used for the LDP TCP session.

protocols mpls-ldp discovery interfaces hello-holdtime

Configures the time interval for which the MPLS LDP link "Hello" adjacency configuration is maintained.

Syntax:
`set protocols mpls-ldp discovery interfaces` hello-holdtime *15..3600*

Syntax:
`delete protocols mpls-ldp discovery interfaces` hello-holdtime *15..3600*

Syntax:
`show protocols mpls-ldp discovery interfaces` hello-holdtime *15..3600*

15..3600
The interval in seconds, which can be from 15 through 3600.

Configuration mode

```
protocols {
    mpls-ldp {
        discovery {
            interfaces  {
                hello-holdtime <15..3600>
            }
        }
    }
}
```

Use this command to configure the time interval in seconds for which the MPLS LDP link "Hello" adjacency configuration is maintained in the absence of link "Hello" messages from the MPLS LDP neighbor. The hello-holdtime value must be at least three times the hello-interval value.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use the `set` form of this command to configure the time interval for which the MPLS LDP link "Hello" adjacency configuration is maintained.

Use the `delete` form of this command to delete the configuration of the time interval for which the MPLS LDP link "Hello" adjacency configuration is maintained.

Use the `show` form of this command to display the configuration of the time interval for which the MPLS LDP link "Hello" adjacency configuration is maintained.

Example: Example

The following example shows how to configure 3000 seconds for which the MPLS LDP link "Hello" adjacency configuration is maintained in the absence of link "Hello" messages from the MPLS LDP neighbor.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces hello-holdtime 3000
```

protocols mpls-ldp discovery interfaces hello-interval

Configures the interval between consecutive MPLS LDP link "Hello" messages.

Syntax:
`set protocols mpls-ldp discovery interfaces` hello-interval *5..1200*

Syntax:
`delete protocols mpls-ldp discovery interfaces` hello-interval *5..1200*

Syntax:
`show protocols mpls-ldp discovery interfaces` hello-interval *5..1200*

5..1200
The interval in seconds, which can be from 5 through 1200.

Configuration mode

```
protocols {
    mpls-ldp {
        discovery {
            interfaces  {
                hello-interval <5..1200>
            }
        }
    }
}
```

Use this command to configure the interval in seconds between consecutive MPLS LDP link "Hello" messages used in basic LDP discovery. The hello-holdtime value must be at least three times the hello-interval value.

Use the `set` form of this command to configure the interval between consecutive MPLS LDP link "Hello" messages.

Use the `delete` form of this command to delete the configuration of the interval between consecutive MPLS LDP link "Hello" messages.

Use the `show` form of this command to display the configuration of the interval between consecutive MPLS LDP link "Hello" messages.

Example: Example

The following example shows how to configure 900 seconds between consecutive MPLS LDP link "Hello" messages.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces hello-interval 900
```

protocols mpls-ldp discovery interfaces interface address-family ipv4

Enables MPLS LDP IPv4 operation on the specified interface.

Syntax:
`set protocols mpls-ldp discovery interfaces` interface *name* address-family ipv4

Syntax:
`delete protocols mpls-ldp discovery interfaces` interface *name* address-family ipv4

Syntax:
`show protocols mpls-ldp discovery interfaces` interface *name* address-family ipv4

name

The interface name.

Configuration mode

```
protocols {
    mpls-ldp {
        discovery {
            interfaces  {
                interface <name> {
                    address-family {
                        ipv4
                    }
                }
            }
        }
    }
}
```

Use this command to enable MPLS LDP IPv4 operation on the specified interface.

Use the `set` form of this command to enable LDP on the specified interface.

Use the `delete` form of this command to delete the configuration to enable LDP on the specified interface.

Use the `show` form of this command to display the configuration of LDP on the specified interface.

Example: Example

The following example shows how to enable MPLS LDP IPv4 operation on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p256p1 address-family
 ipv4
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

protocols mpls-ldp discovery interfaces interface hello-holdtime

Sets the time taken before rejecting a peer adjacency on the specified interface.

Syntax:
`set protocols mpls-ldp discovery interfaces` interface *name* hello-holdtime *15..3600*

Syntax:
`delete protocols mpls-ldp discovery interfaces` interface *name* hello-holdtime *15..3600*

Syntax:
`show protocols mpls-ldp discovery interfaces` interface *name* hello-holdtime *15..3600*

15..3600
The hold time interval in seconds, which can be from 15 through 3600.

Configuration mode

```
protocols {
    mpls-ldp {
        discovery {
            interfaces  {
                interface <name> {
                    hello-holdtime <15..3600>
                }
            }
        }
    }
}
```

Use this command to set the hello hold time interval in seconds, the time taken before rejecting a peer adjacency on the specified interface. The hello-holdtime value must be at least three times the hello-interval value.

Use the `set` form of this command to configure the hold time interval.

Use the `delete` form of this command to delete the configuration of the hold time interval.

Use the `show` form of this command to display the configuration of the hold time interval.

Example: Example

The following example shows how to configure 3000 seconds for the hold time interval for the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p256p1 hello-holdtime
 3000
```

protocols mpls-ldp discovery interfaces interface hello-interval

Configures the interval between consecutive MPLS LDP link "Hello" messages on the specified interface.

Syntax:
`set protocols mpls-ldp discovery interfaces` interface *name* hello-interval *5..1200*

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

`delete protocols mpls-ldp discovery interfaces` interface *name* hello-interval *5..1200*

Syntax:
`show protocols mpls-ldp discovery interfaces` interface *name* hello-interval *5..1200*

5..1200
The interval in seconds, which can be from 5 through 1200.

Configuration mode

```
protocols {
    mpls-ldp {
        discovery {
            interfaces  {
                interface <name> {
                    hello-interval <5..1200>
                }
            }
        }
    }
}
```

Use this command to configure the interval in seconds between consecutive MPLS LDP link "Hello" messages sent from the specified interface to peers used in basic LDP discovery. The hello-holdtime value must be at least three times the hello-interval value.

Use the `set` form of this command to set the interval between consecutive MPLS LDP link "Hello" messages.

Use the `delete` form of this command to delete the configuration of the interval between consecutive MPLS LDP link "Hello" messages.

Use the `show` form of this command to display the configuration of the interval between consecutive MPLS LDP link "Hello" messages.

Example: Example

The following example shows how to configure 900 seconds between consecutive MPLS LDP link "Hello" messages from the interface named dp0p1s15 to its peers.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s15 hello-interval 900
```

protocols mpls-ldp discovery interfaces interface igp-synchronization-delay

Sets the interval that the MPLS LDP LSR waits before notifying the IGP that label exchange is completed.

Syntax:
`set protocols mpls-ldp discovery interfaces` interface *name* igp-synchronization-delay *3..60*

Syntax:
`delete protocols mpls-ldp discovery interfaces` interface *name* igp-synchronization-delay *3..60*

Syntax:
`show protocols mpls-ldp discovery interfaces` interface *name* igp-synchronization-delay *3..60*

name
The interface name.

3..60
The delay in seconds, which can be from 3 through 60.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Configuration mode

```
protocols {
    mpls-ldp {
        discovery {
            interfaces  {
                interface <name> {
                    address-family {
                        igp-synchronization-delay <3..60>
                    }
                }
            }
        }
    }
}
```

Use this command to set the interval in seconds that the MPLS LDP LSR waits before notifying the Interior Gateway Protocol (IGP) that label exchange is completed so that IGP can start advertising the normal metric for the link.

Use the `set` form of this command to set the interval that the LDP waits before notifying the Interior Gateway Protocol (IGP) that label exchange is completed.

Use the `delete` form of this command to delete the configuration to set the interval that the LDP waits before notifying the Interior Gateway Protocol (IGP) that label exchange is completed.

Use the `show` form of this command to display the configuration of the interval that the LDP waits before notifying the Interior Gateway Protocol (IGP) that label exchange is completed.

Example: Example

The following example shows how to set 40 seconds for the interval that the MPLS LDP LSR waits before notifying the Interior Gateway Protocol (IGP) that label exchange is completed on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p256p1 igp-
synchronization-delay 40
```

protocols mpls-ldp discovery interfaces interface session-ka-holdtime

Specifies the session keep alive hold time in seconds.

Syntax:
`set protocols mpls-ldp discovery interfaces` interface *name* session-ka-holdtime *30..3600*

Syntax:
`delete protocols mpls-ldp discovery interfaces` interface *name* session-ka-holdtime *30..3600*

Syntax:
`show protocols mpls-ldp discovery interfaces` interface *name* session-ka-holdtime *30..3600*

30..3600
The keep alive hold time interval in seconds, which can be from 30 through 3600.

Configuration mode

```
protocols {
    mpls-ldp {
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
        discovery  {
            interfaces {
                interface <name> {
                    session-ka-holdtime <30..3600>
                }
            }
        }
    }
}
```

Use this command to configure the session keep alive hold time in seconds. The keep alive hold time is the time that the LSR is configured to wait for "keep-alive" messages from the MPLS LDP peers. An inactive LDP session terminates and the corresponding TCP session closes after the specified time.

Use the `set` form of this command to specify the session keep alive hold time.

Use the `delete` form of this command to delete the configuration of the session keep alive hold time.

Use the `show` form of this command to display the configuration that specifies the session keep alive hold time.

Example: Example

The following example shows how to configure 3000 as the interval after which an inactive LDP session on the interface named dp0p1s15 terminates and the corresponding TCP session closes.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s15 session-ka-holdtime
 3000
```

protocols mpls-ldp discovery interfaces interface session-ka-interval

Specifies the session keep alive interval for an interface.

Syntax:
`set protocols mpls-ldp discovery interfaces` interface *name* session-ka-interval *10..1200*

Syntax:
`delete protocols mpls-ldp discovery interfaces` interface *name* session-ka-interval *10..1200*

Syntax:
`show protocols mpls-ldp discovery interfaces` interface *name* session-ka-interval *10..1200*

10..1200
The interval between successive transmissions of keepalive packets, which can be from 10 through 1200.

Configuration mode

```
protocols {
    mpls-ldp {
        discovery  {
            interfaces  {
                interface <name> {
                    session-ka-interval <10..1200>
                }
            }
        }
    }
}
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use this command to configure the interval between successive transmissions of keepalive packets for a session on a specified interface. Keepalive packets are only sent in the absence of other LDP packets transmitted over the LDP session.

Use the `set` form of this command to configure the interval between successive transmissions of keepalive packets.

Use the `delete` form of this command to delete the configuration of the interval between successive transmissions of keepalive packets.

Use the `show` form of this command to display the configuration of the interval between successive transmissions of keepalive packets.

Example: Example

The following example shows how to configure 900 as the interval between successive transmissions of keepalive packets for the interface named dp0p1s3.

```
vyatta@R1# set protocols mpls-ldp discovery interfaces interface dp0p1s3 session-ka-interval
 900
```

protocols mpls-ldp graceful-restart helper-enable

Enables graceful restart helper mode.

Syntax:

`set protocols mpls-ldp graceful-restart` helper-enable

Syntax:

`delete protocols mpls-ldp graceful-restart` helper-enable

Syntax:

`show protocols mpls-ldp graceful-restart` helper-enable

Configuration mode

```
protocols {
    mpls-ldp {
        graceful-restart {
                helper-enable
        }
    }
}
```

Use this command to enable graceful restart helper mode.

Use the `set` form of this command to enable graceful restart helper mode.

Use the `delete` form of this command to delete the configuration that enables graceful restart helper mode.

Use the `show` form of this command to display the configuration that enables graceful restart helper mode.

protocols mpls-ldp graceful-restart reconnect-time

Specifies the time interval that the remote MPLS LDP peer must wait for the local LDP peer to reconnect.

Syntax:

`set protocols mpls-ldp graceful-restart` reconnect-time*10..1800*

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

`delete protocols mpls-ldp graceful-restart` reconnect-time*10..1800*

Syntax:
`show protocols mpls-ldp graceful-restart` reconnect-time*10..1800*

10..1800
The reconnect time, which can be from 10 through 1800 seconds.

Configuration mode

```
protocols {
    mpls-ldp {
        graceful-restart {
                reconnect-time <10..1800>
        }
    }
}
```

Use this command to specify the time interval in seconds that the remote LDP peer must wait for the local LDP peer to reconnect after the remote peer detects the LDP communication failure.

Use the `set` form of this command to to specify the time interval that the remote LDP peer must wait for the local LDP peer to reconnect.

Use the `delete` form of this command to delete the configuration that specifies the time interval that the remote LDP peer must wait for the local LDP peer to reconnect.

Use the `show` form of this command to display the configuration that specifies the time interval that the remote LDP peer must wait for the local LDP peer to reconnect.

Example: Example

The following example shows how to configure 900 seconds as the time interval that the remote LDP peer must wait for the local LDP peer to reconnect after the remote peer detects the LDP communication failure.

```
vyatta@R1# set protocols mpls-ldp graceful-restart reconnect-time 900
```

protocols mpls-ldp graceful-restart recovery-time

Specifies the maximum time that stale label FEC bindings are retained.

Syntax:
`set protocols mpls-ldp graceful-restart` recovery-time*30..3600*

Syntax:
`delete protocols mpls-ldp graceful-restart` recovery-time*30..3600*

Syntax:
`show protocols mpls-ldp graceful-restart` recovery-time*30..3600*

30..3600
The recovery time, which can be from 30 through 3600 seconds.

Configuration mode

```
protocols {
    mpls-ldp {
        graceful-restart {
                recovery-time <30..3600>
        }
    }
}
```

Use this command to specify the maximum time in seconds that stale label FEC bindings are retained.

Use the `set` form of this command to specify the maximum time in seconds that stale label FEC bindings are retained.

Use the `delete` form of this command to delete the configuration that specifies the maximum time in seconds that stale label FEC bindings are retained.

Use the `show` form of this command to display the configuration that specifies the maximum time in seconds that stale label FEC bindings are retained.

Example: Example

The following example shows how to configure 3000 as the interval after which an inactive LDP session terminates and the corresponding TCP session closes for the neighbor.

```
vyatta@R1# set protocols mpls-ldp graceful-restart recovery-time 3000
```

protocols mpls-ldp log dsm

Logs the DSM events.

Syntax:

```
set protocols mpls-ldp log dsm
```

Syntax:

```
delete protocols mpls-ldp log dsm
```

Syntax:

```
show protocols mpls-ldp log dsm
```

Configuration mode

```
protocols {
    mpls-ldp {
        log {
            dsm
        }
    }
}
```

Use this command to log the Downstream State Machine (DSM).

Use the `set` form of this command to log the DSM.

Use the `delete` form of this command to delete the configuration to log the DSM.

Use the `show` form of this command to display the configuration to log the DSM.

protocols mpls-ldp log events

Logs general MPLS LDP events.

Syntax:

```
set protocols mpls-ldp log events
```

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
delete protocols mpls-ldp log events
```

Syntax:
```
show protocols mpls-ldp log events
```

Configuration mode

```
protocols {
    mpls-ldp {
        log {
            events
        }
    }
}
```

Use this command to log general MPLS LDP events.

Use the `set` form of this command to log general LDP events.

Use the `delete` form of this command to delete the configuration to log general LDP events.

Use the `show` form of this command to display the configuration to log general LDP events.

protocols mpls-ldp log fsm

Logs MPLS LDP FSM events.

Syntax:
```
set protocols mpls-ldp log fsm
```

Syntax:
```
delete protocols mpls-ldp log fsm
```

Syntax:
```
show protocols mpls-ldp log fsm
```

Configuration mode

```
protocols {
    mpls-ldp {
        log {
            fsm
        }
    }
}
```

Use this command to log MPLS LDP Finite State Machine (FSM) events.

Use the `set` form of this command to log LDP FSM events.

Use the `delete` form of this command to delete the configuration to log LDP FSM events.

Use the `show` form of this command to display the configuration to log LDP FSM events.

protocols mpls-ldp log nsm

Logs interactions with the NSM.

Syntax:
```
set protocols mpls-ldp log nsm
```

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
delete protocols mpls-ldp log nsm
```

Syntax:

```
show protocols mpls-ldp log nsm
```

Configuration mode

```
protocols {
    mpls-ldp {
        log {
            nsm
        }
    }
}
```

Use this command to log interactions with the Network Services Module (NSM).

Use the `set` form of this command to log interactions with the NSM.

Use the `delete` form of this command to delete the configuration to log interactions with the NSM.

Use the `show` form of this command to display the configuration to log interactions with the NSM.

protocols mpls-ldp log packet address | all | hello | init | keepalive | label | notification

Logs filtered LDP packet events.

Syntax:

```
set protocols mpls-ldp log packet address | all | hello | init | keepalive | label | notification
```

Syntax:

```
delete protocols mpls-ldp log packet address | all | hello | init | keepalive | label | notification
```

Syntax:

```
show protocols mpls-ldp log packet address | all | hello | init | keepalive | label | notification
```

Configuration mode

```
protocols {
    mpls-ldp {
        log {
            packet | all | hello | init | keepalive | label | notification
        }
    }
}
```

Use this command to log LDP packet events filtered by address, hello, init, keepalive, label, and notification messages.

Use the `set` form of this command to log filtered LDP packet events.

Use the `delete` form of this command to delete the configuration to log filtered LDP packet events.

Use the `show` form of this command to display the configuration to log filtered LDP packet events.

protocols mpls-ldp log rib

Logs interactions with the RIB.

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
set protocols mpls-ldp log rib
```

Syntax:

```
delete protocols mpls-ldp log rib
```

Syntax:

```
show protocols mpls-ldp log rib
```

Configuration mode

```
protocols {
    mpls-ldp {
        log {
            rib
        }
    }
}
```

Use this command to log interactions with the Routing Information Base (RIB).

Use the `set` form of this command to log interactions with the RIB.

Use the `delete` form of this command to delete the configuration to log interactions with the RIB.

Use the `show` form of this command to display the configuration to log interactions with the RIB.

protocols mpls-ldp log usm

Logs the USM events.

Syntax:

```
set protocols mpls-ldp log usm
```

Syntax:

```
delete protocols mpls-ldp log usm
```

Syntax:

```
show protocols mpls-ldp log usm
```

Configuration mode

```
protocols {
    mpls-ldp {
        log {
            usm
        }
    }
}
```

Use this command to log the Upstream State Machine (USM).

Use the `set` form of this command to log the USM.

Use the `delete` form of this command to delete the configuration to log the USM.

Use the `show` form of this command to display the configuration to log the USM.

protocols mpls-ldp lsr-id

Sets the local MPLS LDP LSR ID to the specified IP address.

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

`set protocols mpls-ldp lsr-id` *ipaddr*

Syntax:
`delete protocols mpls-ldp lsr-id` *ipaddr*

Syntax:
`show protocols mpls-ldp lsr-id` *ipaddr*

ipaddr
The IP address.

Configuration mode

```
protocols {
    mpls-ldp {
        lsr-id <ipaddr>
    }
}
```

Use this command to set the local MPLS LDP LSR (Label Switch Router) ID to the specified IP address. The LSR-ID must be unique amongst the LDP LSRs. Unless otherwise specified, LDP will make use of the first configured IP address on a loopback interface.

Use the `set` form of this command to associate the local MPLS LDP LSR ID to the specified IP address.

Use the `delete` form of this command to delete the association of the local MPLS LDP LSR ID to the specified IP address.

Use the `show` form of this command to display the association of the local MPLS LDP LSR ID to the specified IP address.

Example: Example

The following example shows how to associate the local MPLS LDP LSR ID to the IP address 192.166.3.2.

```
vyatta@R1# set protocols mpls-ldp lsr-id 192.166.3.2
```

protocols mpls-ldp neighbors neighbor md5-password

Configures the specified MD5 password for the corresponding neighbor.

Syntax:
`set protocols mpls-ldp neighbors` neighbor *ip address* md5-password *text*

Syntax:
`delete protocols mpls-ldp neighbors` neighbor *ip address* md5-password *text*

Syntax:
`show protocols mpls-ldp neighbors` neighbor *ip address* md5-password *text*

ip address
The IP address of the neighbor LSR.

text
The MD5 password.

Configuration mode

```
protocols {
    mpls-ldp {
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
        neighbors  {
            neighbor <ip address> {
                md5-password <text>
            }
        }
    }
}
```

Use this command to configure the specified MD5 password for the corresponding neighbor so that when any mismatching TCP segments that are received are then rejected from the neighbor.

Use the `set` form of this command to configure the specified MD5 password for the corresponding neighbor.

Use the `delete` form of this command to delete the configuration of the specified MD5 password for the corresponding neighbor.

Use the `show` form of this command to display the configuration of the specified MD5 password for the corresponding neighbor.

Example: Example

The following example show how to configure pwtext as the MD5 password for the neighbor at IP address 192.166.3.2.

```
vyatta@R1# set protocols mpls-ldp neighbors neighbor 192.166.3.2 md5-password pwtext
```

protocols mpls-ldp neighbors session-downstream-on-demand

Sends a label upon request from a peer.

Syntax:
`set protocols mpls-ldp neighbors` session-downstream-on-demand

Syntax:
`delete protocols mpls-ldp neighbors` session-downstream-on-demand

Syntax:
`show protocols mpls-ldp neighbors` session-downstream-on-demand

Configuration mode

```
protocols {
    mpls-ldp {
        neighbors  {
            session-downstream-on-demand
        }
    }
}
```

Use this command to send a label upon request from a peer. By default labels are sent without waiting for an MPLS LDP label request message

Use the `set` form of this command to send a label upon request from a peer.

Use the `delete` form of this command to delete the configuration that sends a label upon request from a peer.

Use the `show` form of this command to display the configuration sends a label upon request from a peer.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

protocols mpls-ldp neighbors session-ka-holdtime

Configures the time interval after which an inactive LDP session terminates and the corresponding TCP session closes.

Syntax:
`set protocols mpls-ldp neighbors` session-ka-holdtime*30..3600*

Syntax:
`delete protocols mpls-ldp neighbors` session-ka-holdtime*30..3600*

Syntax:
`show protocols mpls-ldp neighbors` session-ka-holdtime*30..3600*

30..3600
The hold time interval, which can be from 30 through 3600.

Configuration mode

```
protocols {
    mpls-ldp {
        neighbors  {
            session-ka-holdtime <30..3600>
        }
    }
}
```

Use this command to configure the time interval after which an inactive LDP session terminates and the corresponding TCP session closes. Inactivity is defined as not receiving LDP packets from the neighbor.

Use the `set` form of this command to configure the time interval after which an inactive LDP session terminates and the corresponding TCP session closes.

Use the `delete` form of this command to delete the configuration of the time interval after which an inactive LDP session terminates and the corresponding TCP session closes.

Use the `show` form of this command to display the configuration of the time interval after which an inactive LDP session terminates and the corresponding TCP session closes.

Example: Example

The following example shows how to configure 3000 as the interval after which an inactive LDP session terminates and the corresponding TCP session closes for the neighbor.

```
vyatta@R1# set protocols mpls-ldp neighbors session-ka-holdtime 3000
```

protocols mpls-ldp neighbors session-ka-interval

Configures the interval between successive transmissions of keepalive packets.

Syntax:
`set protocols mpls-ldp neighbors` session-ka-interval*10..1200*

Syntax:
`delete protocols mpls-ldp neighbors` session-ka-interval*10..1200*

Syntax:
`show protocols mpls-ldp neighbors` session-ka-interval*10..1200*

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

10..1200

The interval, which can be from 10 through 1200, between successive transmissions of keepalive packets.

Configuration mode

```
protocols {
    mpls-ldp {
        neighbors  {
            session-ka-interval <10..1200>
        }
    }
}
```

Use this command to configure the interval between successive transmissions of keepalive packets. Keepalive packets are only sent in the absence of other LDP packets transmitted over the LDP session.

Use the `set` form of this command to configure the interval between successive transmissions of keepalive packets.

Use the `delete` form of this command to delete the configuration of the interval between successive transmissions of keepalive packets.

Use the `show` form of this command to display the configuration of the interval between successive transmissions of keepalive packets.

Example: Example

The following example shows how to configure 900 as the interval between successive transmissions of keepalive packets.

```
vyatta@R1# set protocols mpls-ldp neighbors session-ka-interval 900
```

reset mpls ldp

Resets MPLS LDP sessions.

Syntax:

`reset mpls ldp` [adjacency all | *X.X.X.X*] | [session all | *X.X.X.X*]

adjacency all

All sessions on adjacent LSRs.

X.X.X.X

The address for the session on a specific adjacent LSR.

session all

All sessions.

X.X.X.X

The session on a specific LSR.

Operational mode

Use this command to reset the MPLS LDP session or optionally reset:

- All sessions on adjacent LSRs or a session on a specific LSR.
- All sessions or a session on a specific LSR.

show mpls ldp

Displays the global MPLS LDP information.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:

```
show mpls ldp
```

Operational mode

Use this command to display the global MPLS LDP information.

The following example shows how to display global MPLS LDP information.

```
vyatta@vyatta:~$ show mpls ldp
Router ID                 : 1.1.1.1
LDP Version               : 1
Global Merge Capability   : Merge Capable
Label Advertisement Mode  : Downstream Unsolicited
Label Retention Mode      : Liberal
Label Control Mode        : Independent
Instance Loop Detection   : Off
Request Retry             : Off
Propagate Release         : Disabled
Graceful Restart          : Disabled
Hello Interval            : 5
Targeted Hello Interval   : 15
Hold time                 : 15
Targeted Hold time        : 45
Keepalive Interval        : 10
Keepalive Timeout         : 30
Request retry Timeout     : 5
Transport Address data    :
  Labelspace 0            : 40.1.1.3 (in use)
Import BGP routes         : No
```

show mpls ldp adjacency

Displays all MPLS LDP adjacencies for the LSR.

Syntax:

```
show mpls ldp adjacency
```

Operational mode

Use this command to display all the MPLS LDP adjacencies for the LSR.

The following example shows how to display all the MPLS LDP adjacencies for the LSR.

```
vyatta@vyatta:~$ show mpls ldp adjacency
IP Address                    Intf Name        Holdtime    LDP-Identifier
192.166.1.1                   dp0p1s3          15          1.1.1.1:0
192.166.3.2                   dp0p1s1          15          3.3.3.3:0
```

show mpls ldp advertise-labels

Displays the MPLS LDP label advertisement policy for the LSR.

Syntax:

```
show mpls ldp advertise-labels
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Operational mode

Use this command to display the MPLS LDP label advertisement policy for the LSR.

The following example shows how to display the MPLS LDP label advertisement policy for the LSR.

```
vyatta@vyatta:~$ show mpls ldp advertise-labels
Advertisement spec:
    Enable the distribution of all assigned labels
```

The following example shows how to display the MPLS LDP label advertisement policy for the LSR that has a prefix-list named advertise-prefixes applied by using the `set protocols mpls-ldp address-family ipv4 label-policy advertise prefix-list advertise-prefixes` command.

```
vyatta@vyatta:~$ show mpls ldp advertise-labels
Advertisement spec:
    Prefix list = advertise-prefixes; Peer acl = any peers
```

show mpls ldp discovery

Displays which interfaces are MPLS LDP-enabled or the LDP parameters configured on a specified interface.

Syntax:

`show mpls ldp discovery` [*interface*]

interface

The name of the interface.

Operational mode

Use this command to display all interfaces that are configured for MPLS LDP discovery.

The following example shows how to display which interfaces are MPLS LDP-enabled.

```
 vyatta@vyatta:~$ show mpls ldp discovery
Interface        LDP Identifier           LDP Enabled Version Merge Capability
lo               2.2.2.2:0                Disabled            N/A
dp0s4            2.2.2.2:0                Disabled            N/A
dp0p1s1          2.2.2.2:0                Enabled     IPv4    Merge capable
dp0p1s2          2.2.2.2:0                Enabled     IPv4    Merge capable
dp0p1s3          2.2.2.2:0                Enabled     IPv4    Merge capable
```

The following example shows how to display which LDP parameters are configured on a specified MPLS LDP-enabled interface.

```
vyatta@vyatta:~$ show mpls ldp interface dp0p1s1
Status                     : Enabled
Version                    : IPv4
Primary IP Address         : 192.166.3.1
Interface Type             : Ethernet
Label Merge Capability     : Merge Capable
Hold Time                  : 15
Hello Interval             : 5
Targeted Hello Interval    : 15
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
Targeted Hold Time      : 45
Keepalive Interval      : 10
Keepalive Timeout       : 30
Advertisement Mode      : Downstream Unsolicited
Label Retention Mode    : Liberal
Multicast Hellos        : Enabled
Max PDU Length          : 4096
```

show mpls ldp downstream

Displays all MPLS LDP downstream sessions and exchanged labels.

Syntax:

`show mpls ldp downstream`

Operational mode

Use this command to display all the MPLS LDP downstream sessions and exchanged labels for the LSR.

The following example shows how to display all MPLS LDP downstream sessions and exchanged labels for the LSR.

```
vyatta@vyatta:~$ show mpls ldp downstream
Session peer 1.1.1.1:
  FEC                         Nexthop Addr       State            Label   Req.ID   Attr
  192.168.252.0/24             connected      Established       impl-null  0
  192.166.8.0/30               connected      Established           52491  0
  192.166.4.0/30               connected      Established           52488  0
  192.166.3.0/30               connected      Established           52487  0
  192.166.2.0/30               connected      Established       impl-null  0
  192.166.1.0/30               connected      Established       impl-null  0
  100.2.2.0/24                   invalid      Established           52482  0
  100.2.1.0/24                   invalid      Established           52481  0
  30.1.0.0/24                    invalid      Established       impl-null  0
  10.0.0.0/8                   connected      Established           52480  0
  3.3.3.3/32                   connected      Established           52483  0
  1.1.1.1/32                 192.166.1.1      Established       impl-null  0
Session peer 3.3.3.3:
  FEC                         Nexthop Addr       State            Label   Req.ID   Attr
  192.168.252.0/24             connected      Established       impl-null  0
  192.166.9.0/30               connected             Idle            none  0
  192.166.8.0/30             192.166.3.2      Established       impl-null  0
  192.166.5.0/30               connected             Idle            none  0
  192.166.4.0/30             192.166.3.2      Established       impl-null  0
  192.166.3.0/30               connected      Established       impl-null  0
  192.166.2.0/30             192.166.3.2      Established       impl-null  0
  10.0.0.0/8                   connected      Established           52484  0
  3.3.3.3/32                 192.166.3.2      Established       impl-null  0
  1.1.1.1/32                 192.166.3.2      Established           52490  0
```

show mpls ldp fec

Displays all FEC information.

Syntax:

`show mpls ldp fec`

Operational mode

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use this command to display all FEC (Forwarding Equivalent Class) information known to the LSR.

The following example shows how to display all FEC information known to the LSR.

```
vyatta@vyatta:~$ show mpls ldp fec
LSR codes    : E/N - LSR is egress/non-egress for this FEC,
               L - LSR received a label for this FEC,
               > - LSR will use this route for the FEC
FEC                     Code    Session         Out Label       Nexthop Addr
1.1.1.1/32              NL>     1.1.1.1         impl-null       192.166.1.1
                        NL>     3.3.3.3         52490           192.166.3.2
2.2.2.2/32              E >     non-existent    none            connected
3.3.3.3/32              NL>     1.1.1.1         52483           invalid
                        NL>     3.3.3.3         impl-null       192.166.3.2
10.0.0.0/8              NL      1.1.1.1         52480           invalid
                        NL      3.3.3.3         52484           invalid
                        E >     non-existent    none            invalid
                        E >     non-existent    none            192.168.252.253
30.1.0.0/24             NL      1.1.1.1         impl-null       connected
100.2.1.0/24            NL      1.1.1.1         52481           connected
100.2.2.0/24            NL      1.1.1.1         52482           connected
192.166.1.0/30          NL      1.1.1.1         impl-null       invalid
                        E >     non-existent    none            connected
192.166.2.0/30          NL>     1.1.1.1         impl-null       invalid
                        NL>     3.3.3.3         impl-null       192.166.3.2
192.166.3.0/30          NL      1.1.1.1         52487           invalid
                        NL      3.3.3.3         impl-null       invalid
                        E >     non-existent    none            connected
192.166.4.0/30          NL>     1.1.1.1         52488           invalid
                        NL>     3.3.3.3         impl-null       192.166.3.2
192.166.5.0/30          N       3.3.3.3         none            invalid
                        E >     non-existent    none            connected
192.166.8.0/30          NL>     1.1.1.1         52491           invalid
                        NL>     3.3.3.3         impl-null       192.166.3.2
192.166.9.0/30          N       3.3.3.3         none            invalid
                        E >     non-existent    none            connected
192.168.252.0/24        NL      1.1.1.1         impl-null       invalid
                        NL      3.3.3.3         impl-null       invalid
                        E >     non-existent    none            connected
```

show mpls ldp graceful-restart

Displays the MPLS LDP GR operational state.

Syntax:

`show mpls ldp graceful-restart`

Operational mode

Use this command to display the GR (Graceful Restart) operational state known to the LSR.

The following example shows how to display all peers configured with GR that are known to the LSR.

```
vyatta@vyatta:~$show mpls ldp graceful-restart
Peer IP Address           IF Name    Restart    My State     Timer Value
10.0.2.13                 dp0s5      Incapable  OPERATIONAL  0(No Timers Running)
11.0.1.11                 dp0s4      Capable    HELPER_MODE  101(Re-connect Time)
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

show mpls ldp igp sync

Displays the MPLS LDP interfaces configured for IGP synchronization.

Syntax:

`show mpls ldp igp sync`

Operational mode

Use this command to display the MPLS LDP interfaces configured for IGP synchronization. See also `show ip ospf interface` for LDP-OSPF sync configuration, where the initial cost of 65535 is used until IGP synchronization is achieved.

The following example shows how to display the MPLS LDP interfaces configured for IGP synchronization.

```
vyatta@vyatta:~$ show mpls ldp igp sync
lo is up, line protocol is up
  LDP not configured; LDP-IGP Synchronization not enabled.
dp0s4 is up, line protocol is up
  LDP not configured; LDP-IGP Synchronization not enabled.
dp0p1s1 is up, line protocol is up
  LDP configured; LDP-IGP Synchronization enabled.
  Session IP Address : 3.3.3.3
    Sync status: Achieved
    Delay timer: Configured, 40 seconds, Not Running
dp0p1s2 is up, line protocol is up
  LDP configured; LDP-IGP Synchronization enabled.
  Session IP Address : NONE
    Sync status: Not achieved
    Delay timer: Configured, 40 seconds, Not Running
```

show mpls ldp interface

Displays all interfaces that are enabled for LDP label switching.

Syntax:

`show mpls ldp interface` [*interface*]

interface

The name of the interface.

Operational mode

Use this command to display all interfaces that are enabled for LDP label switching or to display whether a specified interface is enabled for MPLS LDP label switching.

The following example shows how to display which interfaces are enabled for MPLS LDP label switching.

```
vyatta@vyatta:~$ show mpls ldp interface
Interface        LDP Identifier           LDP Enabled Version Merge Capability
lo               2.2.2.2:0                Disabled            N/A
dp0s4            2.2.2.2:0                Disabled            N/A
dp0p1s1          2.2.2.2:0                Enabled     IPv4    Merge capable
dp0p1s2          2.2.2.2:0                Enabled     IPv4    Merge capable
dp0p1s3          2.2.2.2:0                Enabled     IPv4    Merge capable
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

show mpls ldp lsp

Displays the MPLS LDP LSPs.

Syntax:

`show mpls ldp lsp` [detail | host | prefix]

detail

Displays detailed information on the LSPs.

host

Displays the LSP information per host LSP.

prefix

Displays the LSP information per LSP prefix.

Operational mode

Use this command to display the MPLS LDP LSPs (Label Switch Paths) for a given FEC.

The following example shows how to display which interfaces are enabled for MPLS LDP label switching.

```
 vyatta@vyatta:~$ show mpls ldp lsp
FEC IPV4:1.1.1.1/32 ->
  1.1.1.1/32                 192.166.1.1      Established       impl-null  0
  1.1.1.1/32                 192.166.3.2      Established           52490  0
FEC IPV4:2.2.2.2/32 ->
  2.2.2.2/32                   connected      Established            none  0 None
    2.2.2.2/32                Established       impl-null        0     None
    2.2.2.2/32                Established       impl-null        0     None
FEC IPV4:3.3.3.3/32 ->
  3.3.3.3/32                   connected      Established           52483  0
    3.3.3.3/32                Established           52481        0     None
  3.3.3.3/32                 192.166.3.2      Established       impl-null  0
    3.3.3.3/32                Established           52481        0     None
FEC IPV4:10.0.0.0/8 ->
  10.0.0.0/8                   connected      Established           52480  0
    10.0.0.0/8                Established           52484        0     None
    10.0.0.0/8                Established           52484        0     None
  10.0.0.0/8                   connected      Established           52484  0
    10.0.0.0/8                Established           52484        0     None
    10.0.0.0/8                Established           52484        0     None
  10.0.0.0/8             192.168.252.253      Established            none  0 None
    10.0.0.0/8                Established           52484        0     None
    10.0.0.0/8                Established           52484        0     None
FEC IPV4:30.1.0.0/24 ->
  30.1.0.0/24                    invalid      Established       impl-null  0
FEC IPV4:100.2.1.0/24 ->
  100.2.1.0/24                   invalid      Established           52481  0
FEC IPV4:100.2.2.0/24 ->
  100.2.2.0/24                   invalid      Established           52482  0
```

show mpls ldp neighbors

Displays the MPLS LDP neighbors with their respective LDP identifiers.

Syntax:

`show mpls ldp neighbors`

Operational mode

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use this command to display the MPLS LDP neighbors with their LDP identifiers that comprise the lsr-id and label space.

The following example shows how to display the MPLS LDP neighbors.

```
vyatta@vyatta:~$ show mpls ldp neighbors
IP Address                    Intf Name        Holdtime   LDP-Identifier
192.166.1.1                   dp0p1s3          15         1.1.1.1:0
192.166.3.2                   dp0p1s1          15         3.3.3.3:0
```

show mpls ldp routes

Displays the RIB routes known by MPLS LDP.

Syntax:
`show mpls ldp routes`

Operational mode

Use this command to display the RIB routes known by MPLS LDP.

The following example shows how to display the RIB routes.

```
vyatta@vyatta:~$ show mpls ldp routes
Prefix: 0.0.0.0/0   Nexthop: 192.168.252.253   IFINDEX: 8
Prefix: 1.1.1.1/32   Nexthop: 192.166.1.1   IFINDEX: 11
                     Nexthop: 192.166.3.2   IFINDEX: 9
Prefix: 2.2.2.2/32   Nexthop: 0.0.0.0   IFINDEX: 13
Prefix: 3.3.3.3/32   Nexthop: 192.166.3.2   IFINDEX: 9
Prefix: 10.0.0.0/8   Nexthop: 192.168.252.253   IFINDEX: 8
Prefix: 192.166.1.0/30   Nexthop: 0.0.0.0   IFINDEX: 11
Prefix: 192.166.2.0/30   Nexthop: 192.166.3.2   IFINDEX: 9
Prefix: 192.166.3.0/30   Nexthop: 0.0.0.0   IFINDEX: 9
Prefix: 192.166.4.0/30   Nexthop: 192.166.3.2   IFINDEX: 9
Prefix: 192.166.5.0/30   Nexthop: 0.0.0.0   IFINDEX: 10
Prefix: 192.166.8.0/30   Nexthop: 192.166.3.2   IFINDEX: 9
Prefix: 192.166.9.0/30   Nexthop: 0.0.0.0   IFINDEX: 12
Prefix: 192.168.252.0/24   Nexthop: 0.0.0.0   IFINDEX: 8
```

show mpls ldp session

Displays LDP TCP session and state information.

Syntax:
`show mpls ldp session` [*X.X.X.X*]

X.X.X.X
The session IP address.

Operational mode

Use this command to display all LDP TCP session and state information or the detailed session information including the received and send label information.

The following example shows how to display all LDP TCP session and state information.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@vyatta:~$ show mpls ldp session
Peer IP Address            IF Name          My Role    State          KeepAlive
1.1.1.1                    dp0p1s3          Active     OPERATIONAL    30
3.3.3.3                    dp0p1s1          Passive    OPERATIONAL    30
```

show mpls ldp statistics

Displays MPLS LDP packet sent and received statistics.

Syntax:

`show mpls ldp statistics`

Operational mode

Use this command to display the number of sent and received packets per MPLS LDP packet type.

The following example shows the sent and received statistics per MPLS LDP packet type.

```
vyatta@vyatta:~$ show mpls ldp statistics

========================================================
   LSR ID = 2.2.2.2:0 : INTERFACE NAME: dp0p1s2
========================================================
  PacketType                     Total
                            Sent       Received
  Notification                0         1
  Hello                    1223       485
  Initialization              1         1
  Keepalive                 235       235
  Address                     1         1
  Address Withdraw            0         0
  Label Mapping              15        17
  Label Request               0         0
  Label Withdraw              5         5
  Label Release               5         5
  Request Abort               0         0
=======================================================
```

show mpls ldp upstream

Displays all MPLS LDP upstream sessions and exchanged labels.

Syntax:

`show mpls ldp upstream`

Operational mode

Use this command to display all the MPLS LDP upstream sessions and exchanged labels for the LSR.

The following example displays all MPLS LDP upstream sessions and exchanged labels for the LSR.

```
vyatta@vyatta:~$ show mpls ldp upstream
Session peer 1.1.1.1:
  FEC                         State              Label     Req.ID    Attr
  192.166.8.0/30            Established           52482        0     None
  192.166.4.0/30            Established           52480        0     None
  192.168.252.0/24          Established       impl-null        0     None
  192.166.9.0/30            Established       impl-null        0     None
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
  192.166.5.0/30          Established       impl-null        0     None
  192.166.3.0/30          Established       impl-null        0     None
  192.166.2.0/30          Established           52485        0     None
  192.166.1.0/30          Established       impl-null        0     None
  10.0.0.0/8              Established           52484        0     None
  3.3.3.3/32              Established           52481        0     None
  2.2.2.2/32              Established       impl-null        0     None
Session peer 3.3.3.3:
  FEC                       State                 Label     Req.ID    Attr
  192.168.252.0/24        Established       impl-null        0     None
  192.166.9.0/30          Established       impl-null        0     None
  192.166.5.0/30          Established       impl-null        0     None
  192.166.3.0/30          Established       impl-null        0     None
  192.166.1.0/30          Established       impl-null        0     None
  10.0.0.0/8              Established           52484        0     None
  2.2.2.2/32              Established       impl-null        0     None
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

MPLS Traffic Engineering

Using MPLS in traffic engineering

Traffic engineering is the task of routing network traffic to avoid points of congestion and make efficient use of high bandwidth interfaces.

When used as an application of MPLS, traffic engineering involves creating LSPs that make the best use of available network resources; that is, traffic-engineered LSPs. This section explains the process of creating traffic-engineered LSPs.

Creating traffic-engineered LSPs involves the following tasks:

- Gathering information about the network
- Using the gathered information to select optimal paths through the network
- Setting up and maintaining the paths

CSPF calculates a traffic-engineered path

When you configure a signaled Label Switched Path, you specify the address of the egress LER, as well as optional attributes, such as the LSPs priority and bandwidth requirements.

When you enable the signaled LSP, the Constrained Shortest Path First (CSPF) process on the ingress LER uses this information to calculate a traffic-engineered path between the ingress and egress LERs. You can optionally specify a path of LSRs that the LSP must pass through on the way to the egress LER.

CSPF is an advanced form of the Shortest Path First (SPF) process used by IGP routing protocols. The CSPF process on the ingress LER uses the configured attributes of the LSP, user-specified path (when there is one), and the information in the Traffic Engineering Database (TED) to calculate the traffic-engineered path. This process consists of a sequential list of the physical interfaces that packets assigned to this LSP pass through to travel from the ingress LER to the egress LER. The traffic-engineered path takes into account the network topology, available resources, and user-specified constraints. The traffic-engineered path calculated by CSPF may or may not be the same as the shortest path that would normally be calculated by standard IGP routing protocols.

CSPF is enabled by default for signaled LSPs.

Once the path for the LSP has been calculated, RSVP signaling then causes resources to be reserved and labels to be allocated on each LSR specified in the path. This may cause already existing, lower priority LSPs to be preempted. Once resources are reserved on all the LSRs in the path, the signaled LSP is considered to be activated; that is, packets can be forwarded over it.

The following sections provide additional information about the individual components of the process for activating traffic-engineered signaled LSPs.

Signaled LSPs

When the LSP is enabled, RSVP signaling messages travel to each LSR along the calculated path, reserving resources and causing labels to be dynamically associated with interfaces.

Signaled LSPs are configured at the ingress LER only. When a packet is assigned to a signaled LSP, it follows a pre-established path from the LSPs ingress LER to its egress LER. This path can be one of the following:

- A path that traverses an explicitly specified set of MPLS routers.
- The IGP shortest path across the MPLS domain determined from local routing tables.
- A traffic-engineered path calculated by the device using constraints such as bandwidth reservations, administrative groups, and network topology information.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Traffic Engineering Database

An LSR Traffic Engineering Database (TED) stores topology information about the MPLS domain.

The TED for an LSR contains topology information about the nodes in an MPLS domain and the links that connect them.

This topology information is obtained from the OSPF traffic engineering (OSPF-TE) LSAs with traffic engineering extensions. OSPF-TE LSAs with TE extensions have special extensions that contain information about an MPLS enabled interface's traffic engineering metric, bandwidth reservations, and administrative group memberships.

An LSR, when configured to do so, floods OSPF-TE LSAs with TE extensions for its MPLS-enabled interfaces to its neighboring routers in the OSPF area. Other LSRs store the information from the OSPF-TE LSAs with TE extensions in their own Traffic Engineering Databases (TEDs), allowing each LSR in the area to maintain an identical TED describing the MPLS topology. The topology information in the TED is used by the CSPF process when it calculates traffic-engineered paths for signaled LSPs. You can display the contents of the TED for an LSR.

LSP attributes and requirements used for traffic engineering

In addition to the topology information in the TED, CSPF considers attributes and requirements specified in configuration statements for the LSP. The following user-specified parameters are considered when CSPF calculates a traffic-engineered path for a signaled LSP:

# Destination address of the egress LER

# Explicit path to be used by the LSP

# Bandwidth required by the LSP

# Setup priority for the LSP

# Metric for the LSP

# Whether the LSP includes or excludes links belonging to specified administrative groups

Setting traffic engineering parameters for MPLS interfaces

When using constraints to determine a path for an LSP, CSPF takes into account information that is carried in generic opaque LSAs.

This information can be used to set up a path for a new LSP or to preempt an existing LSP so that an LSP with a higher priority can be established.

OSPF-TE LSAs with TE extensions include Type/Length/Value triplets (TLVs) containing the following information:

- Link type (either point-to-point or multiaccess network) (OSPF-TE LSAs only)
- Link ID (for point-to-point links, this is the Router ID of the LSR at the other end of the link; for multi-access links, this is the address of the network's designated router) (OSPF-TE LSAs only)
- IP address of the local interface
- IP address of the remote interface (must exist with point-to-point links)
- Traffic engineering metric for the link
- Maximum bandwidth on the interface
- Maximum reservable bandwidth on the interface
- Unreserved bandwidth on the interface
- Administrative groups to which the interface belongs

Optionally, you can specify the maximum amount of bandwidth that can be reserved on an interface. In addition, you can assign interfaces to administrative groups.

Reserving bandwidth on an interface

OSPF-TE LSAs with TE extensions contain three TLVs related to bandwidth reservation:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

- The Maximum Bandwidth TLV indicates the maximum outbound bandwidth that can be used on the interface. Maximum Bandwidth is the operating speed of the port. This reflects the actual physical bandwidth of the interface. You cannot configure this TLV.
- The Maximum Reservable Bandwidth TLV indicates the maximum bandwidth that can be reserved on the interface. The Maximum Reservable Bandwidth must be configured; there is no default.
- The Unreserved Bandwidth TLV indicates the amount of bandwidth not yet reserved on the interface. This TLV consists of eight octets, indicating the amount of unreserved bandwidth (in kilobits per second) at each of eight priority levels. The octets correspond to the bandwidth that can be reserved with a hold priority of 0 through 7, arranged in increasing order, with priority 0 occurring at the start of the TLV, and priority 7 at the end of the TLV. The value in each of the octets is less than or equal to the maximum reservable bandwidth. The Unreserved Bandwidth TLV itself is not user-configurable, although it is affected by modifications to the reservable bandwidth on an interface, as well as changes to LSPs.

You can configure the maximum reservable bandwidth as an absolute value. It is mandatory to configure the maximum reservable bandwidth if you want to use bandwidth reservation functionality.

Reservable bandwidth configuration considerations

The `reservable-bandwidth` command is configurable on an MPLS-enabled interface at any time. The configuration of the command takes effect immediately upon preemption of the LSP.

When LSP preemption occurs, when the reservable bandwidth required for a specific LSP is not supported on the interface, then the LSP immediately goes down. When this occurs, an IGP advertisement of this configuration change is triggered and flooded throughout all ports on the network because the maximum reservable bandwidth configured on the interface is different from the value that was previously configured.

To configure the maximum reservable bandwidth as an absolute value for MPLS LSPs on the interface, complete the following step.

1. Configure the maximum reservable bandwidth in kbps.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s15 bandwidth-
   constraints maximum-reservable 10000k
   ```

   Result: In this example, the maximum reservable bandwidth is configured to 10000 kbps.

Adding interfaces to administrative groups

Administrative groups, also known as resource classes or link colors, allows you to assign MPLS-enabled interfaces to various classes.

You can place individual interfaces into administrative groups. For example, you can define a group named gold and assign high-bandwidth interfaces to it. When a device calculates the path for an LSP, it can take into account the administrative group to which a interface belongs. You can configure up to 32 administrative groups. By default, an interface does not belong to any administrative groups.

Administrative groups are in the range from 0 through 31. You can set an administrative group by name and number. To set an administrative group by name, first create a name for the group and associate the name with an administrative group number.

1. To assign the MPLS-enabled interface dp0p1s15 to an administrative group named gold, enter the following command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s15 admin-groups gold
   ```

   Result: In this example, the administrative group named gold is selected. After you add interfaces to administrative groups, you can specify which groups can be included or excluded from LSP calculations.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

How CSPF calculates a traffic-engineered path

Using information in the TED in addition to the attributes and requirements of the LSP, CSPF calculates a traffic-engineered path for the LSP by performing the tasks listed below.

1. When more than one LSP needs to be enabled, CSPF selects the LSP for path calculation based on the LSPs setup priority and bandwidth requirement.
2. Eliminate unsuitable links from consideration.

   Info:

   The device examines the topology information in its TED and uses this information to eliminate links from consideration for the traffic-engineered path. A link is eliminated when any of the following are true:

   Choose from:

   - The link does not have enough reservable bandwidth to fulfill the LSPs configured requirements.
   - The LSP has an **include** statement, and the link does not belong to an administrative group in the statement.
   - The LSP has an **exclude** statement, and the link belongs to an administrative group specified in the exclude statement.

3. Using the remaining links, calculate the shortest path through the MPLS domain.

   Info:

   Using the links that were not eliminated in the previous step, the device calculates the shortest path between the ingress and egress LERs. When the LSP is configured to use an explicit path, the device individually calculates the shortest path between each node in the path.

   Info:

   By default, the path calculated by CSPF can consist of no more than 255 hops, including the ingress and egress LERs. You can optionally change this maximum to a lower number.

4. When multiple paths have the same cost, select one of them.

   Info:

   The output of the CSPF process is a traffic-engineered path, a sequential list of the physical interfaces that packets assigned to this LSP pass through to reach the egress LER. Once the traffic-engineered path has been determined, RSVP signaling attempts to establish the LSP on each LSR in the path.

Configuring CSPF interface constraint

Under the default condition, hops configured as interface addresses in an LSP path are resolved to the router ID.

An LSP can be configured that does or does not traverse a specified interface. The `protocols mpls-rsvp tunnels tunnel <name> primary path-selection affinities constraints exclude-any` command forces the CSPF calculation to exclude any interface that was assigned to the excluded administrative group when creating an LSP.

To configure the device to always include a specified interface when forming an LSP, complete the following step.

1. Create a list of affinities to include for the path. The path will only be able to use interfaces bound to the included administrative groups.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection
    affinities constraints include-any
   ```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

How RSVP establishes a signaled LSP

The traffic-engineered path calculated by CSPF consists of a sequential list of physical interface addresses, corresponding to a path from the ingress LER to the egress LER. Using this traffic-engineered path, RSVP establishes the forwarding state and resource reservations on each LSR in the path.

As with OSPF, special extensions for traffic engineering are defined for RSVP. These extensions include the EXPLICIT_ROUTE, LABEL_REQUEST, LABEL, and RECORD_ROUTE objects. These extensions are described in RFC 3209.

The following diagram illustrates how RSVP establishes a signaled LSP.

Figure 7: How RSVP establishes a signaled LSP

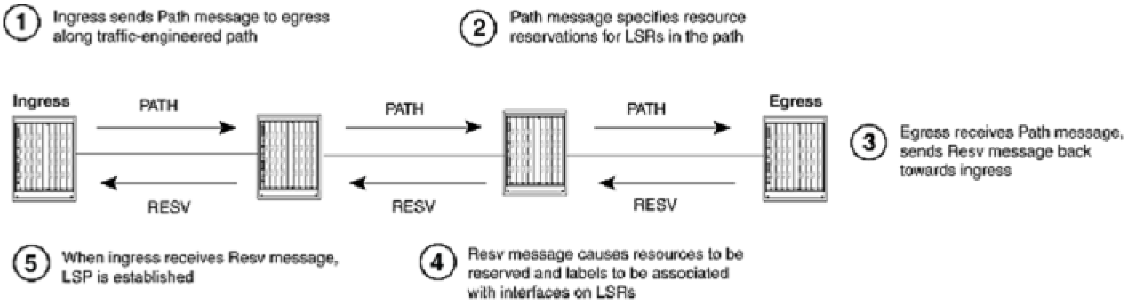
1 Ingress sends Path message to egress along traffic-engineered path
2 Path message specifies resource reservations for LSRs in the path
Ingress
PATH
PATH
PATH
Egress
RESV
RESV
RESV
3 Egress receives Path message, sends Resv message back towards ingress
5 When ingress receives Resv message, LSP is established
4 Resv message causes resources to be reserved and labels to be associated with interfaces on LSRs

RSVP signaling for LSPs works as described below.

1. The ingress LER sends an RSVP Path message towards the egress LER.
   The Path message also describes the traffic for which resources are being requested and specifies the bandwidth that needs to be reserved to accommodate this traffic. In addition, the Path message includes a LABEL_REQUEST object, which requests that labels be allocated on LSRs and tells the egress LER to place a LABEL object in the Resv message that it sends back to the ingress LER.

   Before sending the Path message, the ingress LSR performs admission control on the outbound interface, ensuring that enough bandwidth can be reserved on the interface to meet the LSPs requirements. Admission control examines the LSPs configured setup priority and mean-rate settings. For the LSP to pass admission control, the outbound interface must have reservable bandwidth at the LSPs setup priority level that is greater than the amount of bandwidth specified by the LSPs meanrate setting.
2. The Path message requests resource reservations on the LSRs along the path specified in the ERO.
   When the LSP passes admission control, the ingress LER sends a Path message to the address at the top of the ERO list. This is the address of a physical interface on the next LSR in the path. As the ingress LER did, this LSR performs admission control to make sure the outbound interface has enough reservable bandwidth to accommodate the LSP.

   When the LSP passes admission control, the LSR then removes its address from the top of the ERO list and sends the Path message to the address now at the top of the ERO list. This process repeats until the Path message reaches the last node in the ERO list, which is the egress LER.
3. The egress LER receives the Path message and sends a Resv message towards the ingress LER.
   Resv messages flow upstream from the receiver of the Path message to the sender (that is, from the egress LER to the ingress LER), taking the exact reverse of the path specified in the ERO. In response to the LABEL_REQUEST object in the Path message, the Resv message from the egress LER includes a LABEL object. The LABEL object is used to associate labels with interfaces on the LSRs that make up the LSP.
4. As the Resv messages travel upstream, resources are reserved on each LSR.
   When an LSR receives a Resv message, it again performs admission control on the interface where the Resv message was received (that is, the interface that is the outbound interface for packets traveling through the LSP). When the LSP still passes admission control, bandwidth is allocated to the LSP. The LSR allocates the amount of bandwidth specified by the LSPs meanrate setting, using the bandwidth available to its hold priority level. This may cause lower priority LSPs active on the device to be preempted.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Once bandwidth has been allocated to the LSP, the LABEL object in the Resv message is used to associate labels with interfaces in the LSRs MPLS forwarding table. The following diagram shows an example of how this works.

Figure 8: How the RSVP LABEL object associates a label with an interface in the MPLS forwarding table

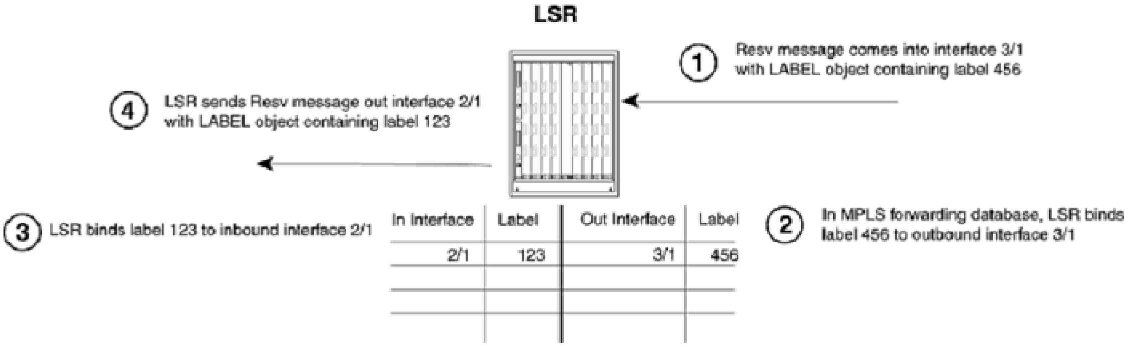

1 Resv message comes into interface 3/1 with LABEL object containing label 456

2 In MPLS forwarding database, LSR binds label 456 to outbound interface 3/1

3 LSR binds label 123 to inbound interface 2/1

4 LSR sends Resv message out interface 2/1 with LABEL object containing label 123

In Interface	Label	Out Interface	Label
2/1	123	3/1	456

In the example above, the LSR receives a Resv message on interface 3/1 from the downstream LSR in the ERO. The Resv message has a LABEL object containing label 456. After performing admission control and bandwidth allocation, the LSR adds an entry to its MPLS forwarding table for this LSP, associating label 456 with outbound interface 3/1.

The LSR then takes a label from its range of available labels (for example, 123) and places it in the LABEL object in the Resv message that it sends to the upstream LSR. In this example, the LSR sends the Resv message out interface 2/1 to the upstream LSR in the ERO. In its MPLS forwarding table for this LSP, the LSR associates label 123 with inbound interface 2/1.

This process repeats at each LSR until the Resv message reaches the ingress LER.

> **Note:** To enable penultimate hop popping for the LSP, the LABEL object sent by the egress LER to the penultimate LSR contains a value of three (3) (Implicit Null Label). This is an IETF-reserved label value that indicates to the penultimate LSR that it must pop the label of MPLS-encoded packets that belong to this LSP.

5. Once the Resv message reaches the ingress LER, and the process described in Step 4 takes place, the LSP is activated. At this point, each LSR in the LSP has reserved resources, allocated labels, and associated labels with interfaces. The LSP is activated, and the ingress LER can assign packets to the LSP.

Refresh messages

Once a signaled LSP is enabled at the ingress LER, the router persistently attempts to establish the LSP through periodic retries until the LSP is successfully established. To maintain the forwarding states and resource reservations on the routers in an LSP, Path and Resv messages are exchanged between neighboring LSRs at regular intervals. When these refresh messages are not received on the routers in the LSP, the RSVP forwarding states and resource reservations are removed. You can also use refresh reduction to reduce RSVP message bandwidth and improve the dependability of RSVP paths and reservations states.

Admission control, bandwidth allocation, and LSP preemption

When a Resv message is received on an LSR, admission control determines whether the LSP can be established, based on its configured priority. When an LSP passes admission control, bandwidth is allocated to the new LSP, possibly preempting existing LSPs that have lower priority.

An LSPs priority consists of a setup priority and a hold priority. The setup priority is the priority for taking resources; the hold priority is the priority for holding resources. An LSPs setup priority is considered during admission control, and its hold priority is considered when bandwidth is allocated to the LSP. The setup and hold priorities are expressed as numbers between zero (0) (highest priority level) and seven (7) (lowest priority level). An LSPs setup priority must be lower than or equal to its hold priority. You can configure either of these values for an LSP; by default, an LSPs setup priority is seven and its hold priority is zero.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

On an MPLS-enabled interface, a certain amount of bandwidth is allocated for usage by LSPs; this amount can be configured as either the maximum available bandwidth on the interface or a portion. The amount of bandwidth an individual LSP can reserve from this pool of allocated bandwidth depends on two user-configured attributes of the LSP: the LSPs priority and the LSPs mean-rate (the average rate of packets that can go through the LSP). The following conditions also apply:

- For an LSP to pass admission control, the bandwidth available to its setup priority level must be greater than the value specified by its mean-rate.
- When an LSP passes admission control, the bandwidth specified by its mean-rate is allocated to the LSP, using bandwidth available to its hold priority level.
- For the allocation of bandwidth to the new LSP, the system might preempt existing, lower-priority LSPs.

When setting up an LSP, the device performs admission control twice: when the Path message is received and when the Resv message is received. when the LSP passes admission control after the Resv message is received, bandwidth allocation and LSP preemption take place.

The sections that follow include examples of how admission control, bandwidth allocation, and preemption work.

Admission control

Admission control examines the LSPs setup priority and mean-rate settings to determine whether the LSP can be activated.

To pass admission control, the reservable bandwidth available at the LSPs setup priority level must be greater than the value specified by its mean-rate.

For example, when the maximum reservable bandwidth on an interface is 10,000 Kbps and no LSPs are currently active, the amount of reservable bandwidth on the interface for each priority level would be as follows:

Priority	Unreserved Bandwidth
0	10,000
1	10,000
2	10,000
3	10,000
4	10,000
5	10,000
6	10,000
7	10,000
Active LSPs: None	

The LSR receives a Resv message for an LSP that has a configured setup priority of six and a hold priority of three. The mean-rate specified for this LSP is 1,000 Kbps. For priority level 6, up to 10,000 Kbps can be reserved. Because the configured mean-rate for this LSP is only 1,000 Kbps, the new LSP passes admission control.

Bandwidth allocation

Once the LSP passes admission control, bandwidth is allocated to the LSP.

The bandwidth allocation procedure examines the LSPs hold priority and mean-rate settings. The amount of bandwidth specified by the mean-rate is allocated to the LSP, using reservable bandwidth available at the LSPs hold priority level.

In this example, the LSPs hold priority is three and mean-rate is 1,000 Kbps. On this interface, for priority level three, up to 10,000 Kbps can be reserved. The amount of bandwidth specified by the mean-rate (1,000 Kbps) is allocated to the LSP.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

After bandwidth is allocated to this LSP, the amount of unreserved bandwidth on the interface is reduced accordingly. In the example, the reservable bandwidth array for the interface now looks like this:

Priority	Unreserved Bandwidth
0	10,000
1	10,000
2	10,000
3	9,000
4	9,000
5	9,000
6	9,000
7	9,000
Active: Lsp with setup 6, hold 3, mean-rate 1,000	

Given the bandwidth allocation above, when an LSP is established with a setup priority of three and a mean-rate of 9,500 Kbps, it would not pass admission control because only 9,000 Kbps is available at priority 3.

LSP preemption

When there is not enough unallocated bandwidth on an interface to fulfill the requirements of a new LSP that has passed admission control, existing LSPs that have a lower priority may be preempted.

When preemption occurs, bandwidth allocated to lower-priority LSPs is reallocated to the higher-priority LSP. LSP preemption depends on the bandwidth requirements and priority of the new LSP, compared to the bandwidth allocation and priority of already existing LSPs.

In the example above, bandwidth has been allocated to an LSP that has a hold priority of three and a mean-rate of 1,000 Kbps. When a new LSP with a setup priority of two, hold priority of one, and mean-rate of 10,000 Kbps is established, admission control, bandwidth allocation, and LSP preemption work as described below.

1. **Admission control:** On the interface, there is 10,000 Kbps available to priority two. The mean-rate for the new LSP is 10,000, so the LSP passes admission control; bandwidth can be allocated to it.
2. **Bandwidth allocation:** The hold priority for the new LSP is one. On the interface, 10,000 Kbps is available to priority one. This entire amount is allocated to the LSP.
3. **LSP preemption:** The first LSP had been using 1,000 Kbps of this amount, but its hold priority is only three. Consequently, the first LSP is preempted, and its bandwidth allocation removed in order to make room for the new LSP.

Once this happens, the reservable bandwidth array for the interface looks like this:

Priority	Unreserved Bandwidth
0	10,000
1	0
2	0
3	0
4	0
5	0
6	0
7	0
Active: LSP with setup 2, hold 1, mean-rate 1,000 Preempted: LSP with setup 6, hold 3, mean-rate 1,000	

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

On this interface, the only LSP that could preempt the active LSP would be have a setup and hold priority of zero.

Enabling OSPF-TE LSAs for MPLS interfaces

Information related to traffic engineering is carried in OSPF traffic engineering (OSPF-TE) LSAs.

OSPF-TE LSAs have special extensions that contain information about an interface's traffic engineering metric, bandwidth reservations, and administrative group memberships.

When an RSVP-enabled device receives an OSPF-TE LSA, it stores the traffic engineering information in its Traffic Engineering Database (TED). The device uses information in the TED when performing calculations to determine a path for an LSP.

By default, OSPF-TE LSAs are sent out for all of its MPLS-enabled interfaces.

> Because information in the TED is used to make path selections using CSPF and information in the TED comes from OSPF-TE LSAs, you do not need to enable the device to send out OSPF-TE LSAs with TE extensions when you want CSPF to perform constraint-based path selection.

Displaying MPLS and RSVP information

You can display the following information about the MPLS configuration on the device:

- Information about MPLS-enabled interfaces on the device
- Statistics about the MPLS-enabled interfaces
- MPLS summary information
- Contents of the Traffic Engineering Database (TED)
- Status information about signaled LSPs configured on the device
- Information about paths configured on the device
- The label applied at each hop in an LSP
- Contents of the MPLS routing table
- RSVP information, including the status of RSVP-enabled interfaces, session information, and statistics
- Information about OSPF-TE LSAs
- MPLS fast reroute information
- MPLS bypass LSP

Displaying RSVP information

You can display global RSVP information, including the version information, the status of RSVP interfaces, RSVP session information, and RSVP statistics.

To display global RSVP information, including the RSVP version number, as well as the refresh interval and refresh multiple, use the `show mpls rsvp` command.

```
vyatta@vyatta:~$ show mpls rsvp
RSVP Version                 : 1
Process uptime               : 5 minutes
Stagger timer                : Not running
RSVP Refresh Reduction       : Enabled
RSVP Message Acknowledgement  : Disabled
Bundle Send                  : Disabled
NSM Connection               : Up
CSPF Connection IPv4         : Up
CSPF Connection IPv6         : Down
CSPF usage                   : Enabled
Reoptimization               : Disabled
RSVP Refresh Timer           : 30
Keep Multiplier              : 3
Acknowledgement Await Timeout : 10
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
Explicit-Null For Direct Conn : Disabled
Local Protection              : Disabled
Hello Receipt                 : Disabled
Hello Interval                : 2000
Hello Timeout                 : 7000
Loop detection                : Enabled (all interface)
Ingress                       : 5.5.5.5
Ingress                       : N/A (not in use)
Penultimate Hop Popping       : Enabled
Refresh PATH msg parsing      : Enabled
Refresh RESV msg parsing      : Enabled
Detour identification         : Sender-Template
Notification                  : Disabled
```

Displaying RSVP administrative groups

Use the `show mpls rsvp` **`admin-groups`** command to display the configured administrative groups.

```
vyatta@vyatta:~$ show mpls rsvp admin-groups
 Admin group detail:
  Value of 1 associated with admin group 'red'
  Value of 2 associated with admin group 'blue'
```

Displaying the status of RSVP interfaces

Use the `show mpls rsvp` **`interface`** command to display the status of RSVP on devices where it is enabled.

```
vyatta@vyatta:~$ show mpls rsvp interface
Interface     RSVP status   Interface Type
lo            Disabled      N/A
eth0          Disabled      N/A
dp0p1s6       Disabled      N/A
dp0p1s10      Enabled       Ethernet
dp0p1s11      Enabled       Ethernet
dp0p1s15      Enabled       Ethernet
```

Use the `show mpls rsvp` **`interface [<name>]`** command to display the RSVP parameters associated with a specific interface.

```
vyatta@vyatta:~$ show mpls rsvp interface dp0p1s11
Status                        : Enabled
Interface Index               : 10
Refresh Reduction usage       : Enabled
Message Acknowledgement       : Disabled
Bundle Buffer size            : 65532
Current Epoch Value           : 247115164
Primary IPv4 address          : 10.10.11.5
Primary IPv6 address          : fe80::5054:ff:fe00:511
Interface Type                : Ethernet
Administrative Group          : blue
Configured refresh time       : 30
Configured keep multiplier    : 3
Acknowledgement Await Timeout : 10
Hello Receipt                 : Disabled
Hello Interval                : 2000
Hello Timeout                 : 7000
Non IANA Hello exchange       : Disabled
```

Displaying RSVP neighbors

Use the `show mpls rsvp` **`neighbor`** command to display a summary of all RSVP neighbors.

```
vyatta@vyatta:~$ show mpls rsvp neighbor
IP Address          UpStrm LSP DnStrm LSP RefreshReduc Srefresh In  Type      GraceRestart
10.10.15.2          1          1          Enabled      24s          Implicit  Incapable
```

Use the `show mpls rsvp` **`neighbor`** *ip-address-of-neighbor* command to display the list of active LSPs for a specific neighbor.

```
vyatta@vyatta:~$ show mpls rsvp neighbor 10.10.15.2
 Nbr Hello State: Down
 Upstream LSPs: 1, Downstream LSPs: 1
 Neighbor supports Refresh Reduction, next SRefresh transmission in: 19s
 Neighbor does not support Graceful Restart.
  Tunnel ID  LSP ID     Ingress             Egress              Type
  5001       101        5.5.5.5             8.8.8.8             Downstream
  5001       101        8.8.8.8             5.5.5.5             Upstream
```

Displaying the RSVP path configuration

Use the `show mpls rsvp path` command to display the path configuration.

```
vyatta@vyatta:~$ show mpls rsvp path
Path name: p678, id: 2, hop-count: 2 type: mpls
 6.6.6.6 loose
 7.7.7.7 loose

Path name: p238, id: 1, hop-count: 3 type: mpls
 2.2.2.2 loose
 3.3.3.3 loose
 8.8.8.8 loose
```

Use the `show mpls rsvp path [<name>]` command to display explicit path configuration.

```
vyatta@vyatta:~$ show mpls rsvp path p678
Path name: p678, id: 2, hop-count: 2 type: mpls
 6.6.6.6 loose
 7.7.7.7 loose
```

Displaying RSVP session information

To display RSVP session information, use the `show mpls rsvp` **`session`** command.

```
vyatta@vyatta:~$ show mpls rsvp session
Ingress RSVP:
To              From            State Pri Rt Style Labelin Labelout LSPname
8.8.8.8         5.5.5.5         Up    Yes  1  1 SE       -   53120 t8
8.8.8.8         10.10.10.5      Up    No   1  1 SE       -   52480 t8
Total 2 displayed, Up 2, Down 0.

Egress RSVP:
To              From            State Pri Rt Style Labelin Labelout LSPname
5.5.5.5         8.8.8.8         Up    Yes  1  1 SE       3        - t5
Total 1 displayed, Up 1, Down 0.
```

To display a count of the total number of configured, ingress, egress and transit RSVP sessions, use the `show mpls rsvp` **`session count`** command.

```
vyatta@vyatta:~$ show mpls rsvp session count
Total configured: 3, Up 3, Down 0

Total ingress sessions: 2, Up 2, Down 0
Total transit sessions: 0, Up 0, Down 0
Total egress sessions: 1, Up 1, Down 0
```

To display detailed session information, use the `show mpls rsvp` **`session [ egress | ingress | transit ] [ down | up ] detail`** command. The default is to show all sessions. You can limit the output to just egress, ingress or transit sessions and/or sessions that are up or down.

```
vyatta@vyatta:~$ show mpls rsvp session egress detail
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
Egress (Primary)
5.5.5.5
  From: 8.8.8.8, LSPstate: Up, LSPname: t5
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in:        3,  Label out: -
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5001, LSP Id: 101, Ext-Tunnel Id: 8.8.8.8
  Upstream: 10.10.15.2, dp0p1s15
  Path lifetime: 157 seconds (due in 120 seconds)
  Resv refresh: 30 seconds (due in 32975 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups:   Received Explicit Route Detail :
   10.10.15.5/32 strict
  Record route: 10.10.18.8 10.10.3.3 10.10.15.2 <self>
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
```

To display detailed session information for a specific tunnel, use the `show mpls rsvp` **`session name <name>`** **`[ primary | secondary ]`** command. The default is to show all sessions. You can limit the output to just the primary or secondary configured session.

```
vyatta@vyatta:~$ show mpls rsvp session name t8 primary
Ingress (Primary)
8.8.8.8
  From: 5.5.5.5, LSPstate: Up, LSPname: t8
  Ingress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  Reoptimization: Disabled
  IGP-Shortcut: Disabled, LSP metric: 25
  LSP Protection: one-to-one
  Label in: -,  Label out:    53120
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5001, LSP Id: 101, Ext-Tunnel Id: 5.5.5.5
  Downstream: 10.10.15.2, dp0p1s15
  Path refresh: 30 seconds (RR enabled) (due in 28 seconds)
  Resv lifetime: 157 seconds (due in 138 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: p238 (in use)
  Configured Explicit Route Detail :
   2.2.2.2/32 loose
   3.3.3.3/32 loose
   8.8.8.8/32 loose
  Session Explicit Route Detail :
   10.10.15.2/32 strict
   10.10.3.3/32 strict
   10.10.18.8/32 strict
  Record route: <self> 10.10.15.2 10.10.3.3 10.10.18.8
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  Last Recorded Error Code: None
  Last Recorded Error Value: None
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
```

Displaying RSVP statistics

The device constantly gathers RSVP statistics. RSVP statistics are collected from the time RSVP is enabled, as well as from the last time the RSVP statistics counters were cleared.

To display the RSVP statistics, use the following command:

```
vyatta@vyatta:~$ show mpls rsvp statistics
  PacketType              Total
                   Sent      Received
  Path                9            2
  PathErr             0            0
  PathTear            2            0
  Resv FF             0            0
  Resv WF             0            0
  Resv SE             2           10
  Resv Err            0            0
  ResvTear            0            0
  ResvConf            0            0
  Hello               0            0
  Bundle              0            0
  Ack                 0            0
  SRefresh          129          131
  Notify              0            0
```

To clear the packet statistics that are displayed by the `show mpls rsvp statistics` command, use the following command:

```
vyatta@vyatta:~$ clear mpls rsvp statistics
```

This command resets the counters listed under "since last clear" for the `show mpls rsvp interface detail` and `show mpls rsvp statistics` commands.

Displaying sessions using summary-refresh

To display sessions using summary-refresh, use the `show mpls rsvp summary-refresh` command.

```
vyatta@vyatta:~$ show mpls rsvp summary-refresh
Neighbor Addr       Tunnel ID  LSP ID     Ingress             Egress
10.10.11.6          5001       101        10.10.10.5          8.8.8.8
10.10.15.2          5001       101        8.8.8.8             5.5.5.5
10.10.15.2          5001       101        5.5.5.5             8.8.8.8
```

Displaying RSVP tunnels

To display a summary list of RSVP tunnels, use the `show mpls rsvp tunnel` command.

```
vyatta@vyatta:~$ show mpls rsvp tunnel
Trunk Name      Trunk ID  Type  # Sess  Egress Address(es)
N/A             5001      P2P   1       5.5.5.5
t8              5001      P2P   2       8.8.8.8
Total trunks configured: 2.
```

MPLS fast reroute using one-to-one backup

MPLS Fast Reroute provides the ability for an LSP to route traffic around a failed node by using a detour route as described in RFC 4090. By using the one-to-one backup method, each LSR except the egress router is identified as a Point of Local Repair (PLR). Each PLR tries to initiate a detour LSP to provide a backup route for the protected path. This detour LSP is used to reroute traffic locally on the detour path in the event of a failure on the protected path. The detour path is computed to exclude the protected link or protected node.

Link protection for FRR

To avoid loss of traffic, Fast Reroute (FRR) protects the LSP and allows a broken LSP to be repaired immediately at the point of failure.

A Label Switched Path (LSP) set up across a MPLS network is used to switch traffic across MPLS network. The path used by an LSP across the network is based upon network resources or any other traffic engineering constraints provided by you. Based on TE-constraints, the ingress MPLS router computes the path to be taken by LSP and signals it using RSVP protocol.

By nature, nodes and links in a MPLS network are prone to failure. It is likely that the link or the nodes through which LSP is traversing can fail. In the event of a failure of a node or link, RSVP protocol has mechanisms that inform the ingress node about the failure. On receipt of failure message for LSP across the path, the ingress router re-signals the LSP using a new path.

Due to messaging and other network delays, the ingress router cannot respond fast enough to minimalize the loss of traffic. Traffic is lost from the moment the failure occurs and until the new path is setup for the LSP, which is quite large in quantum for service provider networks.

To avoid loss of traffic, Fast Reroute (FRR) protects the LSP and allows a broken LSP to be repaired immediately at the point of failure. The point of failure is termed as "Point of local repair" (PLR), where the LSP can be repaired locally without intimating or waiting for the ingress router. PLR is the MPLS router which detects the failure and redirects the traffic appropriately to its backup path with minimal loss.

Typically at the PLR, two types of protection can be provided to LSP:

Link Protection: In this protection, the backup is selected in such a way that it avoids the failed link which was used earlier by the LSP. Traffic merges back to the main stream from the backup on the very next MPLS router. Refer to following Link protection for FRR illustrating link protection provided at R2 to LSP ingressing from R1 to R4.

Figure 9: Link protection

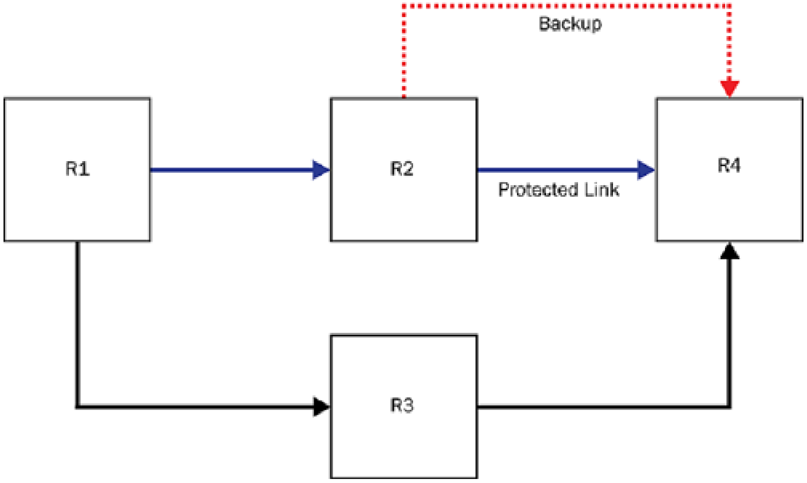

Node Protection: In this protection, backup is selected in such a way that it avoids the failed link along with router to which this link connects. The node which was responsible for link failure is avoided altogether in its entirety, which was used earlier by the LSP. Traffic merges back to main stream from backup on somewhere downstream from the node, which is being avoided. Refer to Link protection for FRR illustrating node protection provided at R1 to LSP ingressing from R1 to R4.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Figure 10: Node protection

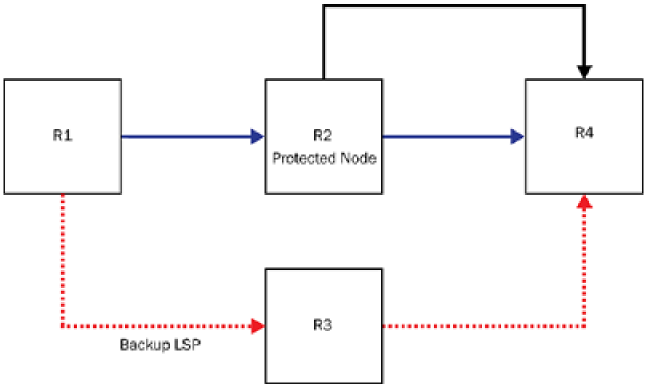

As part of link protection for FRR, ingress routers are allowed to expose this property of MPLS RSVP LSP to you and lets you choose between link protection or node protection. Once the node protection is chosen, PLR first tries to establish a backup LSP, which provides node protection. When node protection is not possible, it attempts to fall back to link protection.

When you choose link protection over node protection, this is communicated to all routers participating in LSP. Each PLR. in this case. limits its search for backup LSP, which provides link protection. In cases where link protection cannot be offered, PLR falls back to node protection.

Link protection for FRR provides options to you to set a preferential method requested for local protection. When RSVP LSP is enabled with FRR (local protection), you are able to configure either link protection or node protection. Link protection is the default.

Path selection metric for CSPF computation

The IGP floods two metrics for every link when the MPLS traffic engineering (TE) is configured in a network. The two metrics are the OSPF link metric and a TE link metric.

The path calculation metric implementation allows you to specify the path calculation for a given tunnel based on either of the following requirements:

- The interior gateway protocol (IGP) link metric for path calculation
- The traffic engineering (TE) link metric for path calculation

The IGP link metric is the default, but you can configure a TE-specific metric on an interface that is used instead for CSPF path computations.

Configuring TE-metric for MPLS interface

How to configure TE-metric for a MPLS interface.

1. MPLS RSVP tunnels must be configured.

   Example:

2. Set the te-metric value at the MPLS interface using the `protocols mpls-rsvp interfaces interface <name> te-metric <1..65535>` command or leave it as a default value to use the igp-metric value of the te-links for CSPF computation (optional).

   Result:

The following example shows how to configure TE-metric for a MPLS interface. In this example, the te-metric is set to 5.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9 te-metric 5
```

Using IGP shortcuts

This feature allows you to configure a signaled LSP to serve as a shortcut between nodes in an AS. In a shortcut LSP, OSPF includes the LSP in the SPF calculation. When OSPF determines that the LSP shortcut is the best path to a destination, it installs a route into the IP routing table, specifying the LSP tunnel interface as the outbound interface, as well as the cost of the LSP. Only LSPs configured to router IDs can be considered as shortcuts. When the LSP goes down, the LSP tunnel route is removed from the main routing table.

The cost of the LSP is the user-configured metric for the LSP. When there is no user-configured metric, the underlying IP cost of the LSP is used. For example, when the IP cost of the best underlying path between two routers is 2, and there is an LSP configured between these two routers, the cost of the LSP is 2. Once an LSP is used as a next hop for a destination, the cost of the LSP can be used to calculate other destinations that can use the LSP egress node as next hop. This allows traffic for addresses downstream from the LSP egress node (including prefixes of the egress node) to use the LSP shortcut.

When OSPF is already using an LSP tunnel route to an Area Border Router (ABR), all inter-area routes through that ABR use the LSP as the next hop, provided there are no other better paths to the destination (paths through other ABRs). An LSP to a destination outside an area is not used by OSPF in the calculation of inter-area routes.

Configuring RSVP-TE

Resource Reservation Protocol - Traffic Engineering (RSVP-TE)

Resource Reservation Protocol - Traffic Engineering (RSVP-TE) is used by applications to reserve resources based on packet stream characteristics.

Enabling MPLS Traffic Engineering

MPLS is enabled for Traffic Engineering (TE) on configured MPLS RSVP interfaces.

MPLS can be used to direct packets through a network over a predetermined path of routers. Traffic engineering works with MPLS to create paths that make the best use of available network resources. In this task, MPLS TE is enabled when RSVP is configured on interfaces.

1. To enable MPLS TE, configure RSVP on the interface and commit the configuration.

   Result:

   ```
   vyatta@vyatta:~$ configure
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p192p1
   vyatta@R1# commit
   vyatta@R1# exit
   vyatta@vyatta:~$
   ```

   The result is that the interface is enabled for MPLS TE.

2. You can enter the `show mpls rsvp interface` command to display that the interface is enabled for MPLS TE.

   Example:

   Result:

   ```
   vyatta@vyatta:~$ show mpls rsvp interface
   Interface     RSVP status   Interface Type
   dp0p192p1     Enabled       Ethernet
   ```

Configuring MPLS RSVP interfaces

How to configure MPLS RSVP interfaces

RSVP-TE must be enabled on interfaces before MPLS RSVP is supported.

To configure MPLS RSVP interfaces, perform one or more of the following steps in configuration mode.

1. To enable RSVP-TE on the interface, enter the `protocols mpls-rsvp interfaces interface <name>` command. Tunnels can only be established over interfaces that are enabled for RSVP-TE.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9
   ```

2. To associate an administrative group with an interface, enter the `protocols mpls-rsvp interfaces interface <name> admin-groups <name>` command. The administrative group must be defined in the globals section. This allows path affinities to exclude or include interfaces associated with specific administrative groups.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9 admin-groups group8
   ```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

3. To set the maximum reservable bandwidth for an interface, enter the `protocols mpls-rsvp interfaces interface <name> bandwidth-constraints maximum-reservable <1-10000000000>` command. This is required to support tunnels with specific bandwidth requirements. The value can be alternatively configured with a k (kilo), m (mega) or g (giga) suffix, for example 1g.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9 bandwidth-constraints
    maximum-reservable 2g
   ```

4. To configure the interval between successive hello packets in milliseconds, enter the `protocols mpls-rsvp interfaces interface <name> signaling hello interval <10..65535>` command. The default interval is 2 seconds. Hello packets are only sent to explicitly configured neighbors. The exchange of hello packets can be use to detect link failures in the absence of a physical notification.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9 signaling hello
    interval 100
   ```

5. To enable the reception of hello packets from a neighbor, enter the `protocols mpls-rsvp interfaces interface <name> signaling hello receipt` command. Incoming hello packets will be ignored if this is not enabled.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9 signaling hello receipt
   ```

6. To configure how long to wait in milliseconds before assuming the link to be dead, enter the `protocols mpls-rsvp interfaces interface <name> signaling hello timeout <10..65535>` command. The default is 7 seconds.

   Example:

   ```
   vyatta@R1#  set protocols mpls-rsvp interfaces interface dp0p1s9 signaling hello
    timeout 350
   ```

7. To configure an interval in seconds for refresh reduction transmissions, enter the `protocols mpls-rsvp interfaces interface <name> signaling refresh interval <10..65535>` command. The default is 30 seconds.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9 signaling refresh
    interval 60
   ```

8. To disable refresh reduction procedures, enter the `protocols mpls-rsvp interfaces interface <name> signaling refresh reduction disable` command. The default is enabled.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9 signaling refresh
    reduction disable
   ```

9. To configure the interface metric to be used with TE tunnels, enter the `protocols mpls-rsvp interfaces interface <name> te-metric <1..65535>` command. By default, the IGP metric is used.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p1s9 te-metric <1..65535>
   ```

Establishing administrative group names

Administrative groups, also known as resource classes or link colors, allow you to assign MPLS-enabled interfaces to various classes.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

When a device calculates the path for an LSP, it can take into account the administrative group to which an interface belongs; you can specify which administrative groups the device can include or exclude when making its calculation.

As many as 32 administrative groups can be configured on the device. You can see an administrative group either by its name or its number. Before you can see an administrative group by its name, you must specify a name for the group at the MPLS policy level and associate the name with the number of that administrative group.

To establish an administrative group name, perform the following steps.

1. Enable the name of the global administrative group.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals name gold value 30
   ```

2. Configure the name of the administrative group for the interface.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp interfaces interface <name> admin-groups gold 30
   ```

   Result: In this example, the administrative group name gold is used with the administrative group number 30. The administrative group number ranges from 0 through 31.

After you associate an administrative group name with a number, you can see it by name when assigning interfaces to the group or including or excluding the group from LSP calculations.

Configuring MPLS RSVP global settings

How to configure MPLS RSVP global settings

Global MPLS RSVP settings can be configured for administrative groups and explicit path characteristics.

To configure global MPLS RSVP settings, perform one or more of the following steps in configuration mode.

1. To create a name to value binding for an administrative group, enter the `protocols mpls-rsvp globals admin-groups <name> value <0-31>` command. This can be referenced by an interface or a path affinity.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals admin-groups group8 value 1
   ```

2. To create an explicit path definition, enter the `protocols mpls-rsvp globals explicit-paths <name>` command. This can be referenced by a tunnel primary or secondary session. It describes a full or partial path that the session must take. Changes to the path will result in make-before-break sessions being re-established with the new information.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals explicit-paths <name>
   ```

3. To create an explicit route object for a path at a position indicated by the index value, enter the `protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects <0-255>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects
    8
   ```

4. To create an explicit route object that is a loose next hop, enter the `protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects <0-255> action loose` command. The explicit route object is a loose next hop - it doesn't have to immediately follow the preceding hop.

   Example:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@R1# set protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects
 8 action loose
```

5. To create an explicit route object that is a strict next hop, enter the `protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects <0-255> action strict` command. The explicit route object is a strict next hop - it must immediately follow the preceding hop.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects
    8 action strict
   ```

6. To configure the address of the next hop in the path, enter the `protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects <0-255> address <x.x.x.x>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects
    8 address 10.10.10.10
   ```

7. To enable periodic reoptimization of tunnels, enter the `protocols mpls-rsvp globals reoptimization` command. The default reoptimization time is 3600 seconds. If reoptimization is not enabled, a session will remain on its current path even if a better path becomes available.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals reoptimization
   ```

8. To configure a specific reoptimization interval in seconds, enter the `protocols mpls-rsvp globals reoptimization interval <0..604800>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals reoptimization interval 30
   ```

9. To configure use of an explicit-null label at the tail end, enter the `protocols mpls-rsvp globals tail-signaling explicit-null` command. The default is to use implicit-null.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals tail-signaling explicit-null
   ```

Resetting RSVP sessions

Reset RSVP sessions

RSVP sessions must be preset.

The default is to tear down and re-establish all sessions. Optionally just primary or secondary sessions.

To reset MPLS RSVP sessions, perform one or more of the following steps.

1. To reset sessions for all tunnels, enter the `reset mpls rsvp tunnel all [ primary | secondary ]` command.

   Example:

   ```
   vyatta@R1# reset mpls rsvp tunnel all
   ```

2. To reset ingress sessions for either all tunnels or a specific tunnel, enter the `reset mpls rsvp tunnel ingress all | name <tunnel-name>` command.

   Example:

   ```
   vyatta@R1# reset mpls rsvp tunnel ingress all
   ```

3. To reset sessions for a specific tunnel, enter the `reset mpls rsvp tunnel name <name> [ primary | secondary ]` command. The default is to tear down and re-establish all sessions. Optionally just primary or secondary sessions.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example:

```
vyatta@R1# reset mpls rsvp tunnel name t7
```

4. To reset non-ingress sessions for either all tunnels or a specific tunnel, enter the `reset mpls rsvp tunnel non-ingress all | name <tunnel-name>` command.

   Example:

   ```
   vyatta@R1# reset mpls rsvp tunnel non-ingress all
   ```

5. To re-optimize the sessions for either all tunnels or a specific tunnel, enter the `reset mpls rsvp tunnel reoptimize all | name <tunnel-name> [ primary | secondary ]` command. The default is to re-optimize all sessions. Optionally just primary or secondary sessions.

   Example:

   Reoptimization involves establishing a new make-before-break session if there is a better path available. No new session is established if there is no better path.

   ```
   vyatta@R1# reset mpls rsvp tunnel reoptimize all t7
   ```

Configuring MPLS RSVP neighbors

How to configure MPLS RSVP neighbors to exchange hello packets

RSVP-TE must be enabled on interfaces before MPLS RSVP is supported.

You can configure a neighbor to exchange hello packets in order to detect link failures.

1. To configure a neighbor to exchange hello packets, enter the `protocols mpls-rsvp neighbors neighbor <x.x.x.x>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp neighbors neighbor 10.10.10.10
   ```

Configuring MPLS RSVP tunnels

How to configure MPLS RSVP tunnels

RSVP-TE must be enabled on interfaces before TE is supported.

To configure RVSP-TE tunnels (signaled LSPs) perform these steps:

1. To configure the head-end for an RSVP-TE tunnel, enter the `protocols mpls-rsvp tunnels tunnel <name>` command. No sessions will be established unless a destination is configured.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2
   ```

2. To enable calculation of routes over TE-tunnels by IGPs, enter the `protocols mpls-rsvp tunnels tunnel <name> autoroute-announce` command. By default, only the route to the tunnel destination is added to the RIB. When enabled, IGPs will treat a tunnel as a single link when calculating routes.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 autoroute-announce
   ```

3. To configure a fixed metric to be used for the tunnel by the IGP for its shortest path calculations, enter the `protocols mpls-rsvp tunnels tunnel <name> autoroute-announce absolute-metric <1..65535>` command. The default is the IGP metric. This command is mutually exclusive with the `protocols mpls-rsvp tunnels tunnel <name> autoroute-announce relative-metric <-65535..65535>` command.

   Example:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 autoroute-announce absolute-
metric  <1..65535>
```

4. To configure a metric that is relative to the IGP metric, enter the `protocols mpls-rsvp tunnels tunnel <name> autoroute-announce relative-metric <-65535..65535>` command. This command is mutually exclusive with the `protocols mpls-rsvp tunnels tunnel <name> autoroute-announce absolute-metric <-65535..65535>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 autoroute-announce relative-
   metric <-65535..65535>
   ```

5. To configure the address of the tail end of the tunnel, enter the `protocols mpls-rsvp tunnels tunnel <name> destination <x.x.x.x>` command. This is required in order to establish any sessions.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 destination 4.4.4.4
   ```

6. To configure the source address of the tunnel, enter the `protocols mpls-rsvp tunnels tunnel <name> source <x.x.x.x>` command. This defaults to the address configured on the first interface, usually a loopback interface.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 source 2.2.2.2
   ```

Configuring MPLS RSVP primary path

How to configure MPLS RSVP primary paths

RSVP-TE must be enabled on interfaces before MPLS RSVP is supported and the MPLS RSVP tunnel must be configured.

You can configure several attributes for an RSVP primary path, such as the bandwidth, affinities to include or exclude, and so on.

To configure an MPLS RSVP primary path, perform one or more of the following steps in configuration mode.

1. To configure the bandwidth to be reserved along the primary path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary bandwidth <1-10000000000>` command. The interfaces must be configured with a maximum reservable bandwidth before tunnels can reserve bandwidth on them. By default, no bandwidth is reserved.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary bandwidth 100000000
   ```

2. To configure the primary path to follow the full or partial explicit path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary explicit-path <name>` command. The primary path will follow the full or partial explicit path, which must be defined in the globals section. CSPF calculates the best path between each pair of nodes in the explicit path.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary explicit-path <name>
   ```

3. To configure the primary path so that it cannot be reoptimized, enter the `protocols mpls-rsvp tunnels tunnel <name> primary lockdown` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary lockdown
   ```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

4. To create a list of affinities to include for the path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary path-selection affinities constraints include-any` command. The path will only be able to use interfaces bound to the included administrative groups.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection
    affinities constraints include-any
   ```

5. To configure the name of administrative group to include for the path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary path-selection affinities constraints include-any affinity-names <name>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection
    affinities constraints include-any affinity-names <name>
   ```

6. To create a list of affinities to exclude for the path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary primary path-selection affinities constraints exclude-any` command. The path will not be able to use interfaces bound to the excluded administrative groups.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary  path-selection
    affinities constraints exclude-any
   ```

7. To configure the name of administrative group to exclude for the path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary path-selection affinities constraints exclude-any affinity-names <name>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection
    affinities constraints exclude-any affinity-names <name>
   ```

8. To configure a hold priority for this session , enter the `protocols mpls-rsvp tunnels tunnel <name> primary priority hold <0-7>` command. The default is 0 (highest), which means it will not be pre-empted by higher priority sessions. Priority 7 is the lowest priority.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary priority hold 3
   ```

9. To configure a setup priority for this session, enter the `protocols mpls-rsvp tunnels tunnel <name> primary priority setup <0-7>` command. The default is 7 (lowest), which means it cannot pre-empt lower priority sessions.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary priority setup 3
   ```

Configuring MPLS RSVP fast reroute

How to configure MPLS RSVP fast reroute

RSVP-TE must be enabled on interfaces before MPLS RSVP is supported. The MPLS RSVP tunnel and the primary path must also be configured to support fast reroute.

You can configure fast reroute, along with several primary path attributes, on the primary path.

To configure an MPLS RSVP fast reroute, perform one or more of the following steps in configuration mode.

1. To enable fast-reroute for the primary path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute` command. This will establish node or link protection at each hop along the primary path.

   Example:

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute
```

2. To configure a specified bandwidth for fast-reroute on the primary path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute bandwidth <1-10000000000>` command. The bandwidth to be reserved along the detour.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute
    bandwidth 100000000
   ```

3. To configure node protection along the path for fast-reroute on the primary path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute node-protection-desired` command. The default is link protection.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute node-
   protection-desired
   ```

4. To create a list of affinities to include for the fast-reroute path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute path-selection affinities constraints include-any` command. The path will only be able to use interfaces bound to the included administrative groups.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute path-
   selection affinities constraints include-any
   ```

5. To configure the name of administrative group to include for the fast reroute path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute path-selection affinities constraints include-any affinity-names <name>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute path-
   selection affinities constraints include-any affinity-names <name>
   ```

6. To create a list of affinities to exclude for the fast reroute path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute path-selection affinities constraints exclude-any` command. The path will not be able to use interfaces bound to the excluded administrative groups.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute path-
   selection affinities constraints exclude-any
   ```

7. To configure the name of administrative group to exclude for the fast reroute path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute path-selection affinities constraints exclude-any affinity-names <name>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute path-
   selection affinities constraints exclude-any affinity-names <name>
   ```

8. To configure a hold priority for this fast reroute session , enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute priority hold <0-7>` command. The default is 0 (highest), which means it will not be pre-empted by higher priority sessions. Priority 7 is the lowest priority.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute priority
    hold 3
   ```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

9. To configure a setup priority for this fast reroute session, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute priority setup <0-7>` command. The default is 7 (lowest), which means it cannot pre-empt lower priority sessions.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute priority
    setup 3
   ```

10. To configure protection one-to-one for this fast reroute path, enter the `protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute protection one-to-one` command. This is the default.

    Example:

    ```
    vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute
     protection one-to-one
    ```

Configuring MPLS RSVP secondary path

How to configure MPLS RSVP secondary paths

RSVP-TE must be enabled on interfaces before MPLS RSVP is supported and the MPLS RSVP tunnel must be configured.

You can configure the same attributes for an RSVP secondary path as for the primary path, such as the bandwidth, affinities to include or exclude, and so on, except for fast-reroute.

To configure an MPLS RSVP secondary path, perform one or more of the following steps in configuration mode.

1. To configure the bandwidth to be reserved along the secondary path, enter the `protocols mpls-rsvp tunnels tunnel <name> secondary bandwidth <1-10000000000>` command. The interfaces must be configured with a maximum reservable bandwidth before tunnels can reserve bandwidth on them. By default, no bandwidth is reserved.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary 100000000
   ```

2. To configure the secondary path to follow the full or partial explicit path, enter the `protocols mpls-rsvp tunnels tunnel <name> secondary explicit-path <name>` command. The secondary path will follow the full or partial explicit path, which must be defined in the globals section. The default is to follow the CSPF calculated path.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary explicit-path
    <name>
   ```

3. To configure the secondary path so that it cannot be reoptimized, enter the `protocols mpls-rsvp tunnels tunnel <name> secondary lockdown` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary lockdown
   ```

4. To create a list of affinities to include for the path, enter the `protocols mpls-rsvp tunnels tunnel <name> secondary path-selection affinities constraints include-any` command. The path will only be able to use interfaces bound to the included administrative groups.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary path-selection
    affinities constraints include-any
   ```

5. To configure the name of administrative group to include for the path, enter the `protocols mpls-rsvp tunnels tunnel <name> secondary path-selection affinities constraints include-any affinity-names <name>` command.

Example:

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary path-selection
 affinities constraints include-any affinity-names <name>
```

6. To create a list of affinities to exclude for the path, enter the `protocols mpls-rsvp tunnels tunnel <name> secondary secondary path-selection affinities constraints exclude-any` command. The path will not be able to use interfaces bound to the excluded administrative groups.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary  path-selection
    affinities constraints exclude-any
   ```

7. To configure the name of administrative group to exclude for the path, enter the `protocols mpls-rsvp tunnels tunnel <name> secondary path-selection affinities constraints exclude-any affinity-names <name>` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary path-selection
    affinities constraints exclude-any affinity-names <name>
   ```

8. To configure a hold priority for this session , enter the `protocols mpls-rsvp tunnels tunnel <name> secondary priority hold <0-7>` command. The default is 0 (highest), which means it will not be pre-empted by higher priority sessions. Priority 7 is the lowest priority.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary priority hold 3
   ```

9. To configure a setup priority for this session, enter the `protocols mpls-rsvp tunnels tunnel <name> secondary priority setup <0-7>` command. The default is 7 (lowest), which means it cannot pre-empt lower priority sessions.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary priority setup 3
   ```

Enabling MPLS RSVP-TE monitoring

How to enable MPLS RSVP-TE monitoring

RSVP-TE must be enabled on interfaces before MPLS RSVP is supported.

You can enable or disable all or specific RSVP-TE event monitoring.

1. To enable RSVP-TE event monitoring, enter the `monitor protocols mpls rsvp enable [ cspf | events | fsm | nsm | packet | rib ]` command.

   Example:

   ```
   vyatta@vyatta:~$ monitor protocols mpls rsvp enable nsm
   ```

2. To disable RSVP-TE event monitoring, enter the `monitor protocols mpls rsvp disable [ cspf | events | fsm | nsm | packet | rib ]` command.

   Example:

   ```
   vyatta@vyatta:~$ monitor protocols mpls rsvp disable nsm
   ```

Configuring MPLS RSVP logging

How to configure MPLS RSVP logging

MPLS RSVP must be enabled.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

You can configure logging to record Constrained Shortest Path First (CSPF) information, general RSVP-TE events, ingress and egress state machine events, downstream and upstream machine events, packet transmission and reception, and so on.

To configure MPLS RSVP logging, perform one or more of the following steps in configuration mode.

1. To configure logging to record CSPF information, enter the `protocols mpls-rsvp log cspf` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log cspf
   ```

2. To configure logging to record general RSVP-TE events, enter the `protocols mpls-rsvp log events` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log events
   ```

3. To configure logging to record egress state machine events, enter the `protocols mpls-rsvp log fsm egress` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log fsm egress
   ```

4. To configure logging to record ingress state machine events, enter the `protocols mpls-rsvp log fsm ingress` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log fsm ingress
   ```

5. To configure logging to record transit downstream state machine events, enter the `protocols mpls-rsvp log fsm transit downstream` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log fsm transit downstream
   ```

6. To configure logging to record transit upstream state machine events, enter the `protocols mpls-rsvp log fsm transit upstream` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log fsm transit upstream
   ```

7. To configure logging to record the interactions with the Network Services Module, enter the `protocols mpls-rsvp log nsm` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log nsm
   ```

8. To configure logging to record RSVP-TE packet transmission and reception events, enter the `protocols mpls-rsvp log packet` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log packet
   ```

9. To configure logging to record interactions with the Routing Information Base, enter the `protocols mpls-rsvp log rib` command.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp log rib
   ```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

RSVP refresh reduction

RSVP control traffic (Path and Resv messages) is initially propagated to establish an RSVP session and reserve resources along the path, or to signal a change of state (trigger messages). However, because it is a soft-state protocol, RSVP also requires periodic refreshing to prevent reserved resources from aging out. The original RSVP as defined in RFC 2205 achieves this by re-sending identical Path and Resv messages (refresh messages) at regular intervals along the reserved path as long as the RSVP session remains unchanged. The bandwidth and processing time required to support these refresh messages increases linearly as more RSVP sessions are established, which can result in scaling problems.

RFC 2961 establishes extensions to RSVP which can help reduce the overhead caused by refresh messages: bundle messages, which allows multiple RSVP messages to be aggregated into a single PDU, and summary refresh messages, which replace identical RSVP message re-transmissions with a list of the IDs of all Path and Resv states to be refreshed.

When you enable either of the refresh reduction extensions on an interface, outgoing RSVP packets sent on that interface sets the refresh reduction capability bit in the common RSVP header to indicate that the AT&T device is capable of receiving and processing refresh reduction messages and related objects.

Setting up signaled LSPs

An LSP consists of an actual path of MPLS routers through a network, as well as the characteristics of the path, including bandwidth allocations and routing metrics.

Signaled LSPs are configured at the ingress LER. When you enable a signaled LSP, RSVP causes resources to be allocated on the other routers in the LSP.

Configuring a signaled LSP consists of the following tasks:

- Specifying a path for the LSP to follow (optional)
- Setting parameters for the signaled LSP
- Specifying which packets are to be forwarded along the LSP (optional)

Setting up paths

A path is a list of router hops that specifies a route across an MPLS domain.

Once you create a path, you can create signaled LSPs that use the path. Paths are configured separately from LSPs so that a path may be specified once and then used by several LSPs that see the path by name. There can be one primary path and one optional secondary path.

A path is always configured at the ingress LER and assumes that the ingress LER is the beginning of the path. A path can contain any number of nodes, which correspond to MPLS-enabled routers in the network. Each node has one attribute: whether it is strict or loose. A strict node means that the router must be directly connected to the preceding node. A loose node means that there can be other routers in between.

Creating a path is not absolutely necessary when configuring an LSP. When you configure a signaled LSP without naming a path, CSPF uses only information in the Traffic Engineering Database (TED), as well as the user-configured attributes and requirements of the LSP to calculate the path.

1. Define the explicit route path.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals explicit-paths sf_to_sj
   ```

   Info: In this example the explicit route path is sf to sj.

2. Configure the loose node.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals explicit-paths <name> explicit-route-
   objects <0-255> action loose
   ```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@R1# set protocols mpls-rsvp globals explicit-paths <name> explicit-route-
objects <0-255> address <ip-address>
```

3. Configure the strict node.

   Example:

   ```
   vyatta@R1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects <0-255> action strict
   vyatta@R1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects <0-255> address <ip-address>
   ```

The path is assumed to start from the local node. You specify the nodes in order from ingress to egress. Specifying the local node itself as the first node in the path is optional. Further, the final node does not necessarily have to be the egress LER in the LSP. (The egress LER is specified at the LSP configuration level with the to command.) When the final node in the path differs from the egress LER, the hop between the final node in the path and the egress LER is treated as a hop to a loose node; that is, standard IP routing is used to determine the path between the final node and the egress LER.

The IP address defines an LSR and can be any interface address or a loopback interface address on the LSR.

The **strict** and **loose** parameters are relative to the preceding node. When specifying a loose node, there can be other routers between the previous node and this one. When specifying a strict node, you must make sure that the LSR is actually directly connected to the preceding node.

Modifying a path

Once you have created a path, you can modify the path.

Complete the following steps to modify a path.

1. For this example, a three node path named *sf_to_sj* was configured as follows.

   Example:

   ```
   vyatta@vm1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 10 address 2.2.2.2
   vyatta@vm1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 20 address 4.4.4.4
   vyatta@vm1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 30 address 6.6.6.6
   ```

2. To modify the path named *sf_to_sj* to include a node between the first and second node, insert a new entry using a number between 10 and 20 for the explicit-route-objects parameter.

   Example:

   ```
   vyatta@vm1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 15 address 3.3.3.3
   ```

3. Modify the *sf_to_sj* path again to delete the node with explicit-route-objects set to 20 (this node became the third node in step 2; it was originally the second node in step 1).

   Example:

   ```
   vyatta@vm1# delete protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 20
   ```

4. Display the resulting *sf_to_sj* path.

   Result:

   ```
   vyatta@vm1# show protocols mpls-rsvp globals explicit-paths
    explicit-paths sf_to_sj {
           explicit-route-objects 10 {
   ```

```
                address 2.2.2.2
        }
        explicit-route-objects 15 {
                address 3.3.3.3
        }
        explicit-route-objects 30 {
                address 6.6.6.6
        }
 }
```

Inserting a node into a path

Once you have created a path, you can add a node into the path.

The <0-255> number in the explicit-route-objects is used to order nodes, so you can insert a new node by using a number in between the numbers used by the two existing nodes. For example, if you have a node number 10 and a node number 20, then you can insert a node by using node number 15. To insert a node into a path complete the following steps.

1. For this example, a three node path named *sf_to_sj* was configured as follows.

   Example:

   ```
   vyatta@vm1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 10 address 2.2.2.2
   vyatta@vm1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 20 address 4.4.4.4
   vyatta@vm1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 30 address 6.6.6.6
   ```

2. To modify the path named *sf_to_sj* to include a node between the first and second node, insert a new entry using a number between 10 and 20 for the explicit-route-objects parameter.

   Example:

   ```
   vyatta@vm1# set protocols mpls-rsvp globals explicit-paths sf_to_sj explicit-route-
   objects 15 address 3.3.3.3
   ```

3. Display the resulting four nodes*sf_to_sj* path.

   Result:

   ```
   vyatta@vm1# show protocols mpls-rsvp globals explicit-paths
    explicit-paths sf_to_sj {
           explicit-route-objects 10 {
                   address 2.2.2.2
           }
           explicit-route-objects 15 {
                   address 3.3.3.3
           }
           explicit-route-objects 20 {
                   address 4.4.4.4
           }        }
           explicit-route-objects 30 {
                   address 6.6.6.6
           }
    }
   ```

> **Note:** When you modify a path, new make-before-break sessions are established using the new path. When successfully established, the new sessions using the new path replace the sessions that used the old path.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Deleting a path

Once you have created a path, you can delete the path.

To delete an entire path from the LSRs configuration, use the delete form of the `protocols mpls-rsvp globals explicit-paths` command as shown for the *sf_to_sj* path.

```
vyatta@vm1# delete protocols mpls-rsvp globals explicit-paths sf_to_sj
```

Configuring MPLS fast reroute using one-to-one backup

To enable FRR with one-to-one backup, complete the following steps:

1. Configure fast-reroute on the primary session for a tunnel named "sf_to_sj":

   Example:

   ```
   vyatta@vm1# set protocols mpls-rsvp tunnels tunnel sf_to_sj primary fast-reroute
   ```

2. In order to use RSVP hello messages to detect link failures and trigger a switch over to the backup session, configure the neighbors to exchange hello messages and enable hello message receipt on the interfaces:

   Example:

   ```
   vyatta@vm1# set protocols mpls-rsvp neighbors neighbor 10.10.100.2
   vyatta@vm1# set protocols mpls-rsvp interfaces interface dp0p1s1 signaling hello
    receipt
   ```

Optionally, you can configure the hello transmission interval and timeout (signaling hello interval | timeout).

```
vyatta@vm1# set protocols mpls-rsvp interfaces interface dp0p1s1 signaling hello interval 100
```

```
vyatta@vm1# set protocols mpls-rsvp interfaces interface dp0p1s1 signaling hello timeout 300
```

MPLS fast reroute using one-to-one backup configuration options

The following options can be set for a MPLS fast reroute path using one-to-one backup configuration:

- Bandwidth
- Node Protection Desired
- Path Selection
- Priority
- Protection

These options are displayed as follows:

```
vyatta@vm1# set protocols mpls-rsvp tunnels tunnel pe1-pe2
 primary fast-reroute
Possible Completions:
   <Enter>                  Execute the current command
   bandwidth                Bandwidth to be reserved (bps) <1-10000000000>
   node-protection-desired Request FRR node protection on LSRs
 > path-selection           Path selection properties
 > priority                 Priority to use in the session attribute object
   protection               LSP Protection mechanism
```

For information on using these configuration options, see Configuring MPLS RSVP Fast Reroute *(page 103)*.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Configuring the graceful restart helper mode

Many routers have separate control plane routing and forwarding plane processes. Therefore, it is possible for a particular routing process to fail independently of the forwarding plane. In such a scenario, it is often desirable for the forwarding plane to be unaffected by the routing process failure and for this process to restart and reconstitute its prefailure control plane state, typically with assistance from its neighbors or peers.

Many protocols, including RSVP-TE (as described in RFC 3209, RFC 3473, and RFC 5063), describe that a graceful-restart procedure works as follows.

1. MPLS neighbors detect, usually by means of a liveliness detection mechanism (for example, Hello messages), the failure and restarting of the control plane process for the RSVP-TE protocol on a neighboring router.
2. Neighbors of the failing router attempt to minimize data-plane disruption (usually by pretending that the neighbor is still up).
3. Neighbors of the failing router attempt to help it recover its control plane state after it restarts, usually by retransmitting the appropriate protocol messages.

A router whose routing process fails goes into the restart mode and the neighbors of the router switch to the helper mode. However, because routing processes cannot restart on the vRouter, the vRouter supports only the RSVP-TE graceful-restart helper mode. The restarting mode is not supported.

When in the RSVP-TE graceful-restart helper mode, a vRouter can help, with respect to the LSP direction, an upstream or a downstream neighbor that is restarting as long as the neighbor indicates that it supports RSVP-TE graceful restart in accordance with the relevant RFCs.

To configure graceful restart on your vRouter, use the following command:

`set protocols mpls-rsvp globals graceful-restart`

For example, to enable the RSVP-TE graceful restart helper mode on your vRouter, enter the following commands:

```
vyatta@vyatta# set protocols mpls-rsvp globals graceful-restart
vyatta@vyatta# commit
```

> **Note:** You must explicitly configure the neighbors with which to exchange the graceful restart capability, as shown in Configuring MPLS RSVP neighbors *(page 101)*.

Disabling CSPF path calculations

By default, CSPF is enabled for signaled LSP calculations, but you can disable CSPF. For example, to allow a TE tunnel LSP to traverse OSPF areas, disable CSPF.

Disabling CSPF means that the full CSPF path is not computed up front. Instead, the vRouter queries the RIB for the next-hop information. This information is needed to reach the egress vRouter. Then, the vRouter sends the RSVP PATH message to the next hop without an Explicit Route Object (ERO). A similar process occurs at each hop along the path.

If you configure a full or partial explicit path, the vRouter queries the RIB for the next-hop information to reach the first hop in the explicit path. The vRouter sends the RSVP PATH message with an ERO whose first hop is the actual first next hop, and whose subsequent hops are from the configured path. Each hop along the path removes the first hop from the ERO and performs a RIB lookup on the next hop in the list.

> **Note:** If you configure constraints for a CSPF path, the next hop returned by the RIB query must meet the configured constraints or the session fails to come up—no alternative path can be found when CSPF is disabled. Each hop along the path checks whether its next hop meets any constraints present in the PATH message.

To disable constraint-based path selection for a TE tunnel, use the following command:

`set protocols mpls-rsvp tunnels tunnel` *name* { **`primary`** | **`secondary`** } **`no-cspf`**

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

For example, to disable constraint-based path selection for the primary and secondary path options for tunnel1, a TE tunnel, enter the following commands:

```
vyatta@vyatta# set protocols mpls-rsvp tunnels tunnel tunnel1 primary no-cspf
vyatta@vyatta# set protocols mpls-rsvp tunnels tunnel tunnel1 secondary no-cspf
vyatta@vyatta# commit
```

To verify that CSPF is disabled, run the following command:

```
vyatta@vyatta:~$ show mpls rsvp session detail
Ingress (Primary)
6.6.6.6
  From: 1.1.1.1, LSPstate: Up, LSPname: tunnel1
  Ingress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  CSPF usage: Disabled
  Reoptimization: Disabled
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: -,  Label out: -
  Tspec rate: 0, Fspec rate: 0
  ...
```

RSVP Sample Configuration

This section presents an example of a typical MPLS configuration.

Minimum MPLS RSVP example

Minimum configuration example for MPLS RSVP

MPLS RSVP must be enabled on interfaces before MPLS RSVP is supported.

This configuration example shows how to configure MPLS RSVP on the Head node with the vyatta@vm-mpls8-5# prompt, the Midpoint node with the vyatta@vm-mpls8-2# prompt, and the Tail node with the vyatta@vm-mpls8-7# prompt in the following topology:

Head --- Midpoint --- Tail

1. Enter the following configuration commands for the Head node at the vyatta@vm-mpls8-5# prompt.

   Example:

   ```
   set interfaces dataplane dp0p1s15 address 10.10.15.5/24
   set interfaces loopback lo address 5.5.5.5/32
   set protocols mpls-rsvp interfaces interface dp0p1s15
   set protocols mpls-rsvp tunnels tunnel t7 destination 7.7.7.7
   set protocols ospf area 0 network 5.5.5.5/32
   set protocols ospf area 0 network 10.10.15.0/24
   commit
   ```

2. Enter the following command at the vyatta@vm-mpls8-5# prompt to verify the Head node configuration.

   Example:

   ```
   vyatta@vm-mpls8-5# run show mpls rsvp session
   Ingress RSVP:
   To              From            State Pri Rt Style Labelin Labelout LSPname
   7.7.7.7         5.5.5.5         Up    Yes  1  1 SE       -   53120 t7
   Total 1 displayed, Up 1, Down 0.
   ```

3. Enter the following configuration commands for the Midpoint node at the vyatta@vm-mpls8-2# prompt.

   Example:

   ```
   set interfaces dataplane dp0p1s15 address 10.10.15.2/24
   set interfaces dataplane dp0p1s16 address 10.10.16.2/24
   set interfaces loopback lo address 2.2.2.2/32
   set protocols mpls-rsvp interfaces interface dp0p1s15
   set protocols mpls-rsvp interfaces interface dp0p1s16
   set protocols ospf area 0 network 2.2.2.2/32
   set protocols ospf area 0 network 10.10.15.0/24
   set protocols ospf area 0 network 10.10.16.0/24
   commit
   ```

4. Enter the following command at the vyatta@vm-mpls8-2# prompt to verify the Midpoint node configuration.

   Example:

   ```
   vyatta@vm-mpls8-2# run show mpls rsvp session
   ```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
Transit RSVP:
To              From            State Pri Rt Style Labelin Labelout LSPname
7.7.7.7         5.5.5.5         Up    Yes  1  1 SE    53120       3 t7
Total 1 displayed, Up 1, Down 0.
```

5. Enter the following configuration commands for the Tail node at the vyatta@vm-mpls8-7# prompt.

   Example:

```
set interfaces dataplane dp0p1s16 address 10.10.16.7/24
set interfaces loopback lo address 7.7.7.7/32
set protocols mpls-rsvp interfaces interface dp0p1s16
set protocols ospf area 0 network 7.7.7.7/32
set protocols ospf area 0 network 10.10.16.0/24
commit
```

6. Enter the following command at the vyatta@vm-mpls8-7# prompt to verify the Tail node configuration.

   Example:

```
vyatta@vm-mpls8-7# run show mpls rsvp session
Egress RSVP:
To              From            State Pri Rt Style Labelin Labelout LSPname
7.7.7.7         5.5.5.5         Up    Yes  1  1 SE       3        - t7
Total 1 displayed, Up 1, Down 0.
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

RSVP-TE Commands

RSVP operational commands

The following clear, reset, and show commands are available with MPLS RSVP:

1. clear mpls rsvp statistics
2. reset mpls rsvp tunnel all [primary | secondary]
3. reset mpls rsvp tunnel ingress all | name <tunnel-name>
4. reset mpls rsvp tunnel name <name> [primary | secondary]
5. reset mpls rsvp tunnel non-ingress all | name <tunnel-name>
6. reset mpls rsvp tunnel reoptimize all | name <tunnel-name> [primary | secondary]
7. show mpls rsvp
8. show mpls rsvp admin-groups
9. show mpls rsvp interface [<name>]
10. show mpls rsvp neighbor [X.X.X.X]
11. show mpls rsvp path [<name>]
12. show mpls rsvp session
13. show mpls rsvp session count
14. show mpls rsvp session [egress | ingress | transit] [down | up] detail
15. show mpls rsvp session name <name> [primary | secondary]
16. show mpls rsvp statistics
17. show mpls rsvp summary-refresh
18. show mpls rsvp tunnel

RSVP configuration commands

The following interface configuration commands are available with MPLS RSVP:

1. protocols mpls-rsvp interfaces interface <name>
2. protocols mpls-rsvp interfaces interface <name> admin-groups <name>
3. protocols mpls-rsvp interfaces interface <name> bandwidth-constraints maximum-reservable <1-10000000000>
4. protocols mpls-rsvp interfaces interface <name> signaling hello interval <10..65535>
5. protocols mpls-rsvp interfaces interface <name> signaling hello receipt
6. protocols mpls-rsvp interfaces interface <name> signaling hello timeout <10..65535>
7. protocols mpls-rsvp interfaces interface <name> signaling refresh interval <10..65535>
8. protocols mpls-rsvp interfaces interface <name> signaling refresh reduction disable
9. protocols mpls-rsvp interfaces interface <name> te-metric <1..65535>

The following logging configuration commands are available with MPLS RSVP:

1. protocols mpls-rsvp log cspf
2. protocols mpls-rsvp log events
3. protocols mpls-rsvp log fsm egress
4. protocols mpls-rsvp log fsm ingress
5. protocols mpls-rsvp log fsm transit downstream
6. protocols mpls-rsvp log fsm transit upstream
7. protocols mpls-rsvp log nsm
8. protocols mpls-rsvp log packet
9. protocols mpls-rsvp log rib

Logging is enabled for RSVP with: monitor protocols mpls rsvp enable | disable [cspf | events | fsm | nsm | packet | rib]

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

The following neighbor configuration command is available with MPLS RSVP:

1. protocols mprotocols mpls-rsvp neighbors neighbor <x.x.x.x>

The following tunnel configuration commands are available with MPLS RSVP:

1. protocols mpls-rsvp tunnels tunnel <name>
2. protocols mpls-rsvp tunnels tunnel <name> autoroute-announce
3. protocols mpls-rsvp tunnels tunnel <name> autoroute-announce absolute-metric <1..65535>
4. protocols mpls-rsvp tunnels tunnel <name> autoroute-announce relative-metric <-65535..65535>
5. protocols mpls-rsvp tunnels tunnel <name> destination <x.x.x.x>
6. protocols mpls-rsvp tunnels tunnel <name> source <x.x.x.x>

The following primary path tunnel configuration commands are available with MPLS RSVP:

1. protocols mpls-rsvp tunnels tunnel <name> primary bandwidth <1-10000000000>
2. protocols mpls-rsvp tunnels tunnel <name> primary explicit-path <name>
3. protocols mpls-rsvp tunnels tunnel <name> primary lockdown
4. protocols mpls-rsvp tunnels tunnel <name> primary path-selection affinities constraints include-any
5. protocols mpls-rsvp tunnels tunnel <name> primary path-selection affinities constraints include-any affinity-names <name>
6. protocols mpls-rsvp tunnels tunnel <name> primary path-selection affinities constraints exclude-any
7. protocols mpls-rsvp tunnels tunnel <name> primary path-selection affinities constraints exclude-any affinity-names <name>
8. protocols mpls-rsvp tunnels tunnel <name> primary path-selection hop-limit <value>
9. protocols mpls-rsvp tunnels tunnel <name> primary priority hold <0-7>
10. protocols mpls-rsvp tunnels tunnel <name> primary priority setup <0-7>
11. protocols mpls-rsvp tunnels tunnel <name> primary record-route record-label

The following fast reroute configuration commands are available with MPLS RSVP:

1. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute
2. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute bandwidth <1-10000000000>
3. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute node-protection-desired
4. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute path-selection affinities constraints include-any
5. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute path-selection affinities constraints exclude-any
6. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute path-selection affinities constraints include-any affinity-names
7. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute path-selection affinities constraints exclude-any affinity-names
8. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute priority hold <0-7>
9. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute priority setup <0-7>
10. protocols mpls-rsvp tunnels tunnel <name> primary fast-reroute protection one-to-one

The following secondary path tunnel configuration command is available with MPLS RSVP:

1. protocols mpls-rsvp tunnels tunnel <name> secondary bandwidth <1-10000000000>
2. protocols mpls-rsvp tunnels tunnel <name> secondary explicit-path <name>
3. protocols mpls-rsvp tunnels tunnel <name> secondary lockdown
4. protocols mpls-rsvp tunnels tunnel <name> secondary path-selection affinities constraints include-any
5. protocols mpls-rsvp tunnels tunnel <name> secondary path-selection affinities constraints include-any affinity-names <name>
6. protocols mpls-rsvp tunnels tunnel <name> secondary path-selection affinities constraints exclude-any
7. protocols mpls-rsvp tunnels tunnel <name> secondary path-selection affinities constraints exclude-any affinity-names <name>
8. protocols mpls-rsvp tunnels tunnel <name> secondary path-selection hop-limit <value>

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

9. protocols mpls-rsvp tunnels tunnel <name> secondary priority hold <0-7>
10. protocols mpls-rsvp tunnels tunnel <name> secondary priority setup <0-7>
11. protocols mpls-rsvp tunnels tunnel <name> secondary record-route record-label

Global MPLS RSVP configuration commands

The following global MPLS RSVP configuration commands are available with MPLS:

1. protocols mpls-rsvp globals explicit-paths <name> admin-groups <name> value <0-31>
2. protocols mpls-rsvp globals explicit-paths <name>
3. protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects <0-255>
4. protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects <0-255> action loose
5. protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects <0-255> action strict
6. protocols mpls-rsvp globals explicit-paths <name> explicit-route-objects <0-255> address <x.x.x.x>
7. protocols mpls-rsvp globals reoptimization
8. protocols mpls-rsvp globals reoptimization interval <0..604800>
9. protocols mpls-rsvp globals tail-signaling explicit-null

clear mpls rsvp statistics

Clears the packet statistics that are displayed by the `show mpls rsvp statistics` command.

Syntax:

```
clear mpls rsvp statistics
```

Operational mode

Use this command to clear the packet statistics that are displayed by the `show mpls rsvp statistics` command.

monitor protocol mpls rsvp

Enables or disables the generation of debug messages that are related to MPLS RSVP logs.

Syntax:

```
monitor protocol mpls rsvp { enable { cspf | events | fsm | nsm | packet | rib } | disable { all |
cspf | events | fsm | nsm | packet | rib } }
```

Operational mode

Use this command to enable or disable the generation of debug messages that are related to MPLS RSVP logs.

The following example command enables MPLS RSVP Finite State Machine monitoring.

```
vyatta@vyatta:~$ monitor protocol mpls rsvp enable fsm
```

As a result, the RSVP debugging status is "on" when displayed with the `show monitoring` command.

```
vyatta@vyatta:~$ show monitoring
------------------------------
  Protocol monitoring status
------------------------------

...

RSVP debugging status:
  ...
  RSVP Egress FSM debugging is on
  RSVP Transit Downstream FSM debugging is on
  RSVP Transit Upstream FSM debugging is on
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
   RSVP Ingress FSM debugging is on
...
```

protocols mpls-rsvp globals admin-groups value

Sets a name to value binding for an administrative group.

Syntax:
`set protocols mpls-rsvp globals admin-groups` *name* value *0-31*

Syntax:
`delete protocols mpls-rsvp globals admin-groups` *name* value *0-31*

Syntax:
`show protocols mpls-rsvp globals admin-groups` *name* value *0-31*

name

The name of the administrative group.

value

The value from one through 31 that is associated with the administrative group name.

Configuration mode

```
protocols {
    mpls-rsvp {
        globals  {
            admin-groups <name> {
                value <0-31>
            }
        }
    }
}
```

Use this command to configure a name to value binding for an administrative group. This value can be referenced by an interface or a path affinity.

Use the `set` form of this command to configure a name to value binding for an administrative group.

Use the `delete` form of this command to delete a name to value binding for an administrative group.

Use the `show` form of this command to display the configuration of a name to value binding for an administrative group.

Example: Example

The following example shows how to associate the administrative group named "group10" with a value of 10.

```
vyatta@R1# set protocols mpls-rsvp globals admin-groups group10 value 10
```

protocols mpls-rsvp globals explicit-paths

Creates an explicit path definition.

Syntax:
`set protocols mpls-rsvp globals` explicit-paths *name*

Syntax:

`delete protocols mpls-rsvp globals` explicit-paths *name*

Syntax:
`show protocols mpls-rsvp globals` explicit-paths *name*

name
A path name.

Configuration mode

```
protocols {
    mpls-rsvp {
        globals  {
            explicit-paths <name>
        }
    }
}
```

Use this command to create an explicit path definition. This explicit path definition can be referenced by a tunnel primary or secondary session. It describes a full or partial path that the session must take. Changes to the path result in make-before-break sessions being re-established with the new information.

Use the `set` form of this command to create an explicit path definition.

Use the `delete` form of this command to delete an explicit path definition.

Use the `show` form of this command to display an explicit path definition.

Example: Example

The following example shows how to create an explicit path definition for the path named onegreen.

```
vyatta@R1# set protocols mpls-rsvp globals explicit-paths onegreen
```

protocols mpls-rsvp globals explicit-paths explicit-route-objects

Creates an explicit route object for a specified path at a position indicated by an index value.

Syntax:
`set protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255*

Syntax:
`delete protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255*

Syntax:
`show protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255*

name
A path name.

0-255
The index value, which ranges from 0 through 255, that indicates a position in the path.

Configuration mode

```
protocols {
    mpls-rsvp {
        globals  {
            explicit-paths <name> {
                explicite-route-objects <0-255>
            }
        }
    }
}
```

Use this command to create an explicit route object for a specified path at a position indicated by the specified index value.

Use the `set` form of this command to create an explicit route object.

Use the `delete` form of this command to delete an explicit route object.

Use the `show` form of this command to display an explicit route object.

Example: Example

This example creates an explicit route object positioned at index value 100 in the path named path1.

```
vyatta@R1# set protocols mpls-rsvp globals explicit-path path1 explicit-route-objects 100
```

protocols mpls-rsvp globals explicit-paths explicit-route-objects action loose

Creates an explicit route object that is a loose next-hop for a specified path at a position indicated by an index value.

Syntax:
`set protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* action loose

Syntax:
`delete protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* action loose

Syntax:
`show protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* action loose

name

A path name.

0-255

The index value, which ranges from 0 through 255, that indicates a position in the path.

Configuration mode

```

protocols {
    mpls-rsvp {
        globals  {
            explicit-paths <name> {
                explicite-route-objects <0-255>  {
                    action loose
                }
            }
        }
    }
}
```

Use this command to create an explicit route object (where the object is a loose next-hop) for a specified path at a position indicated by the specified index value. When the explicit route object is a loose next hop then it does not have to immediately follow the preceding hop.

Use the `set` form of this command to create an explicit route object.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use the `delete` form of this command to delete an explicit route object.

Use the `show` form of this command to display an explicit route object.

Example: Example

The following example shows how to create an explicit route loose next hop object positioned at index value 100 in the path named path1.

```
vyatta@R1# set protocols mpls-rsvp globals explicit-path path1 explicit-route-objects 100
 action loose
```

protocols mpls-rsvp globals explicit-paths explicit-route-objects action strict

Creates an explicit route object that is a strict next-hop for a specified path at a position indicated by an index value.

Syntax:
`set protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* action strict

Syntax:
`delete protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* action strict

Syntax:
`show protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* action strict

name

A path name.

0-255

The index value, which ranges from 0 through 255, that indicates a position in the path.

Configuration mode

```
protocols {
    mpls-rsvp {
        globals  {
            explicit-paths <name> {
                explicite-route-objects <0-255> {
                    action strict
                }
            }
        }
    }
}
```

Use this command to create an explicit route object (where the object is a strict next-hop) for a specified path at a position indicated by the specified index value. When the explicit route object is a strict next-hop then it must immediately follow the preceding hop.

Use the `set` form of this command to create an explicit route object.

Use the `delete` form of this command to delete an explicit route object.

Use the `show` form of this command to display an explicit route object.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to create an explicit route strict next-hop object positioned at index value 100 in the path named path1.

```
vyatta@R1# set protocols mpls-rsvp globals explicit-path path1 explicit-route-objects 100
 action strict
```

protocols mpls-rsvp globals explicit-paths explicit-route-objects address

Creates an explicit route object with a specified next-hop address for a specified path at a position indicated by an index value.

Syntax:
`set protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* address *x.x.x.x*

Syntax:
`delete protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* address *x.x.x.x*

Syntax:
`show protocols mpls-rsvp globals` explicit-paths *name* explicit-route-objects *0-255* address *x.x.x.x*

name
A path name.

0-255
The index value, which ranges from 0 through 255, that indicates a position in the path.

x.x.x.x
The address of the next-hop.

Configuration mode

```
protocols {
    mpls-rsvp {
        globals  {
            explicit-paths <name> {
                explicite-route-objects <0-255>  {
                    address X.X.X.X
                }
            }
        }
    }
}
```

Use this command to create an explicit route object (where the object is a next-hop at the specified address) for a specified path at a position indicated by the specified index value. The explicit route object is configured with the address of the next hop in the path.

Use the `set` form of this command to create an explicit route object.

Use the `delete` form of this command to delete an explicit route object.

Use the `show` form of this command to display an explicit route object.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to create an explicit route object with next-hop address 10.10.10.10 and the object is positioned at index value 100 in the path named path1.

```
vyatta@R1# set protocols mpls-rsvp globals explicit-path path1 explicit-route-objects 100
 address 10.10.10.10
```

protocols mpls-rsvp globals graceful-restart

Enables the RSVP-TE graceful-restart helper mode.

Syntax:
`set protocols mpls-rsvp globals graceful-restart`

Syntax:
`delete protocols mpls-rsvp globals graceful-restart`

Syntax:
`show protocols mpls-rsvp globals graceful-restart`

Configuration mode

```
protocols {
       mpls-rsvp {
              globals {
                     graceful-restart
              }
       }
}
```

Use the `set` form of this command to enable the RSVP-TE graceful-restart helper mode.

Use the `delete` form of this command to disable the RSVP-TE graceful-restart helper mode.

Use the `show` form of this command to display the configuration of the RSVP-TE graceful-restart helper mode.

protocols mpls-rsvp globals reoptimization

Enables periodic reoptimization of tunnels.

Syntax:
`set protocols mpls-rsvp globals reoptimization`

Syntax:
`delete protocols mpls-rsvp globals reoptimization`

Syntax:
`show protocols mpls-rsvp globals reoptimization`

Configuration mode

```
protocols {
    mpls-rsvp {
        globals  {
            reoptimization
        }
    }
}
```

Use this command to enable periodic reoptimization of tunnels. The default reoptimization time is 3600 seconds. If reoptimization is not enabled, a session remains on its current path even if a better path becomes available.

Use the `set` form of this command to enable periodic reoptimization of tunnels.

Use the `delete` form of this command to delete periodic reoptimization of tunnels.

Use the `show` form of this command to display the configuration of periodic reoptimization of tunnels.

protocols mpls-rsvp globals reoptimization interval

Configures periodic reoptimization of tunnels at an interval.

Syntax:
`set protocols mpls-rsvp globals reoptimization` interval *0..604800*

Syntax:
`delete protocols mpls-rsvp globals reoptimization` interval *0..604800*

Syntax:
`show protocols mpls-rsvp globals reoptimization` interval *0..604800*

0..604800
The interval value in seconds. The value can be zero through 604800.

Configuration mode

```
protocols {
    mpls-rsvp {
        globals  {
            reoptimization  {
                interval <0..604800>
            }
        }
    }
}
```

Use this command to configure periodic reoptimization of tunnels with a specific reoptimization interval, in seconds. The default reoptimization time is 3600 seconds. If reoptimization is not enabled, a session will remain on its current path even if a better path becomes available.

Use the `set` form of this command to enable periodic reoptimization of tunnels.

Use the `delete` form of this command to delete periodic reoptimization of tunnels.

Use the `show` form of this command to display the configuration of periodic reoptimization of tunnels.

Example: Example

The following example shows how to configure the reoptimization of tunnels to happen every 1000 seconds.

```
vyatta@R1# set protocols mpls-rsvp globals reoptimization interval 1000
```

protocols mpls-rsvp globals tail-signaling explicit-null

Configures an explicit-null label at the tail end.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`set protocols mpls-rsvp globals tail-signaling explicit-null`

Syntax:
`delete protocols mpls-rsvp globals tail-signaling explicit-null`

Syntax:
`show protocols mpls-rsvp globals tail-signaling explicit-null`

Configuration mode

```
protocols {
    mpls-rsvp {
        globals  {
            tail-signaling explicit-null
        }
    }
}
```

Use this command to set an explicit-null label at the tail end of a tunnel. The default is to use implicit-null.

Use the `set` form of this command to set an explicit-null label at the tail end of a tunnel.

Use the `delete` form of this command to delete an explicit-null label at the tail end of a tunnel.

Use the `show` form of this command to display the configuration of an explicit-null label at the tail end of a tunnel.

protocols mpls-rsvp interfaces interface

Enables RSVP-TE on the specified interface.

Syntax:
`set protocols mpls-rsvp interfaces` interface *name*

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name*

Syntax:
`show protocols mpls-rsvp interfaces` interface *name*

name

The interface name.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
            interface <name>
        }
    }
}
```

Use this command to enable RSVP-TE on the specified interface. Tunnels can only be established over interfaces which are enabled for RSVP-TE.

Use the `set` form of this command to enable RSVP-TE on the specified interface.

Use the `delete` form of this command to delete the configuration of RSVP-TE on the specified interface.

Use the `show` form of this command to show the configuration of RSVP-TE on the specified interface.

Example: Example

The following example shows how to enable RSVP-TE on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p256p1
```

protocols mpls-rsvp interfaces interface admin-groups

Associates an administrative group with an interface.

Syntax:
`set protocols mpls-rsvp interfaces` interface *name* admin-groups *name*

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name* admin-groups *name*

Syntax:
`show protocols mpls-rsvp interfaces` interface *name* admin-groups *name*

interface *name*
The interface name.

admin-groups *name*
The administrative group name.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
            interface <name> {
                admin-groups <name>
            }
        }
    }
}
```

Use this command to associate an administrative group with an interface. The administrative group must be previously defined using the `protocols mpls-rsvp globals admin-groups <name> value <0-31>` command. MPLS RSVP must be enabled on the interface. This allows path affinities to exclude or include interfaces associated with specific administrative groups.

Use the `set` form of this command to associate an administrative group with an interface.

Use the `delete` form of this command to delete the association of an administrative group with an interface.

Use the `show` form of this command to display the association of an administrative group with an interface.

Example: Example

The following example shows how to associate the administrative group named group8 with the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p256p1 admin-groups group8
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

protocols mpls-rsvp interfaces interface bandwidth-constraints maximum-reservable

Sets the maximum reservable bandwidth for an interface.

Syntax:
`set protocols mpls-rsvp interfaces` interface *name* bandwidth-constraints maximum-reservable *1-10000000000*

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name* bandwidth-constraints maximum-reservable *1-10000000000*

Syntax:
`show protocols mpls-rsvp interfaces` interface *name* bandwidth-constraints maximum-reservable *1-10000000000*

name
The name of the interface.

1-10000000000
The maximum bandwidth, which can be from 1 through 10000000000.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
            interface <name> {
                bandwidth-constraints {
                    maximum-reservable <1..10000000000>
                }
            }
        }
    }
}
```

Use this command to set the maximum reservable bandwidth for an interface. This is required to support tunnels with specific bandwidth requirements. The value can be alternatively configured with a k (kilo), m (mega) or g (giga) suffix; for example 1g.

Use the `set` form of this command to set the maximum reservable bandwidth for an interface.

Use the `delete` form of this command to delete the configuration of the maximum reservable bandwidth for an interface.

Use the `show` form of this command to display the configured maximum reservable bandwidth for an interface.

Example: Example

The following example shows how to set the maximum reservable bandwidth to 100000 for the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p256p1 bandwidth-constraints
 maximum-reservable 100000
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

protocols mpls-rsvp interfaces interface signaling hello interval

Configures the interval between successive hello packets in milliseconds.

Syntax:
`set protocols mpls-rsvp interfaces` interface *name* signaling hello interval *10..65535*

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name* signaling hello interval *10..65535*

Syntax:
`show protocols mpls-rsvp interfaces` interface *name* signaling hello interval *10..65535*

interface *name*

The interface name.

signaling hello interval*10..65535*

The interval between successive hello packets in milliseconds, which can be from 10 through 65535.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
            interface <name> {
                signaling {
                    hello {
                        interval <10..65535>
                    }
                }
            }
        }
    }
}
```

Use this command to configure the interval between successive hello packets in milliseconds. The default interval is 2 seconds. Hello packets are only sent to explicitly configured neighbors. The exchange of hello packets can be used to detect link failures in the absence of a physical notification.

Use the `set` form of this command to configure the interval between successive hello packets in milliseconds.

Use the `delete` form of this command to delete the configuration of the interval between successive hello packets in milliseconds.

Use the `show` form of this command to display the configured interval between successive hello packets in milliseconds.

Example: Example

The following example shows how to configure the interval between successive hello packets to 6500 milliseconds on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p256p1 signaling hello interval
 6500
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

protocols mpls-rsvp interfaces interface signaling hello receipt

Enables the reception of hello packets from a neighbor.

Syntax:
`set protocols mpls-rsvp interfaces` interface *name* signaling hello receipt

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name* signaling hello receipt

Syntax:
`show protocols mpls-rsvp interfaces` interface *name* signaling hello receipt

interface *name*

The interface name.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
            interface <name> {
                signaling  {
                    hello  {
                        receipt
                    }
                }
            }
        }
    }
}
```

Use this command to enable the reception of hello packets from a neighbor. Incoming hello packets will be ignored if this is not enabled.

Use the `set` form of this command to enable the reception of hello packets from a neighbor.

Use the `delete` form of this command to delete the configuration that enables the reception of hello packets from a neighbor.

Use the `show` form of this command to display the configuration that enables the reception of hello packets from a neighbor.

Example: Example

The following example shows how to enable the reception of hello packets from a neighbor on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p256p1 signaling hello receipt
```

protocols mpls-rsvp interfaces interface signaling hello timeout

Configures the time to wait before assuming the link to be dead.

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

`set protocols mpls-rsvp interfaces` interface *name* signaling hello timeout *10..65535*

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name* signaling hello timeout *10..65535*

Syntax:
`show protocols mpls-rsvp interfaces` interface *name* signaling hello timeout *10..65535*

interface *name*
The interface name.

signaling hello timeout*10..65535*
The timeout between successive hello packets in milliseconds, which can be from 10 through 65535.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
            interface <name> {
                signaling {
                    hello {
                        timeout <10..65535>
                    }
                }
            }
        }
    }
}
```

Use this command to configure how long to wait in milliseconds before assuming the link to be dead. The default is 7 seconds.

Use the `set` form of this command to configure the timeout between successive hello packets in milliseconds.

Use the `delete` form of this command to delete the configuration of the timeout between successive hello packets in milliseconds.

Use the `show` form of this command to display the configured timeout between successive hello packets in milliseconds.

Example: Example

The following example shows how to configure the timeout between successive hello packets to 1000 milliseconds on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p256p1 signaling hello timeout
 1000
```

protocols mpls-rsvp interfaces interface signaling refresh interval

Configures an interval for refresh reduction transmissions.

Syntax:
`set protocols mpls-rsvp interfaces` interface *name* signaling refresh interval *10..65535*

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name* signaling refresh interval *10..65535*

Syntax:
`show protocols mpls-rsvp interfaces` interface *name* signaling refresh interval *10..65535*

interface *name*
The interface name.

signaling refresh interval *10..65535*
The signaling refresh interval value in seconds, which can be from 10 through 65535.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
            interface <name> {
                signaling {
                    refresh {
                        interval <10..65535>
                    }
                }
            }
        }
    }
}
```

Use this command to configure an interval in seconds for refresh reduction transmissions. The default is 30 seconds.

Use the `set` form of this command to configure an interval for refresh reduction transmissions.

Use the `delete` form of this command to delete the configuration of an interval for refresh reduction transmissions.

Use the `show` form of this command to display the configuration of an interval for refresh reduction transmissions.

Example: Example

The following example shows how to configure a 60 second interval for refresh reduction transmissions on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p256p1 signaling refresh interval
 60
```

protocols mpls-rsvp interfaces interface signaling refresh reduction disable

Disables refresh reduction procedures.

Syntax:
`set protocols mpls-rsvp interfaces` interface *name* signaling refresh reduction disable

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name* signaling refresh reduction disable

Syntax:
`show protocols mpls-rsvp interfaces` interface *name* signaling refresh reduction disable

interface *name*

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

The interface name.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
            interface <name> {
                signaling {
                    refresh {
                        reduction {
                            disable
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to disable refresh reduction procedures. The default is enabled.

Use the `set` form of this command to disable refresh reduction procedures.

Use the `delete` form of this command to delete the configuration that disables refresh reduction procedures.

Use the `show` form of this command to display the configuration that disables refresh reduction procedures.

Example: Example

The following example shows how to disable refresh reduction procedures on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface dp0p256p1 signaling refresh reduction
 disable
```

protocols mpls-rsvp interfaces interface te-metric

Configures the interface metric to be used with TE tunnels.

Syntax:
`set protocols mpls-rsvp interfaces` interface *name* te-metric *1..65535*

Syntax:
`delete protocols mpls-rsvp interfaces` interface *name* te-metric *1..65535*

Syntax:
`show protocols mpls-rsvp interfaces` interface *name* te-metric *1..65535*

interface *name*
The interface name.

te-metric *1..65535*
The value for the interface metric to be used with TE tunnels. The value can be from 1 through 65535.

Configuration mode

```
protocols {
    mpls-rsvp {
        interfaces  {
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
            interface <name> {
                te-metric <1..65535>
            }
        }
    }
}
```

Use this command to configure the interface metric to be used with TE tunnels. By default, the IGP metric is used.

Use the `set` form of this command to configure the interface metric to be used with TE tunnels.

Use the `delete` form of this command to delete the configuration of an interface metric to be used with TE tunnels.

Use the `show` form of this command to display the configured interface metric to be used with TE tunnels.

Example: Example

The following example shows how to configure the interface metric 655 to be used with TE tunnels on the interface named dp0p256p1.

```
vyatta@R1# set protocols mpls-rsvp interfaces interface te-metric 655
```

protocols mpls-rsvp log cspf

Logs CSPF information.

Syntax:
`set protocols mpls-rsvp log cspf`

Syntax:
`delete protocols mpls-rsvp log cspf`

Syntax:
`show protocols mpls-rsvp log cspf`

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            cspf
        }
    }
}
```

Use this command to log Constrained Shortest Path First (CSPF) information.

Use the `set` form of this command to log CSPF information.

Use the `delete` form of this command to delete the configuration to log CSPF information.

Use the `show` form of this command to display the configuration to log CSPF information.

protocols mpls-rsvp log events

Logs general RSVP-TE events.

Syntax:
`set protocols mpls-rsvp log events`

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:

```
delete protocols mpls-rsvp log events
```

Syntax:

```
show protocols mpls-rsvp log events
```

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            events
        }
    }
}
```

Use this command to log general RSVP-TE events.

Use the `set` form of this command to log general RSVP-TE events.

Use the `delete` form of this command to delete the configuration to log general RSVP-TE events.

Use the `show` form of this command to display the configuration to log general RSVP-TE events.

protocols mpls-rsvp log fsm egress

Logs egress state machine events.

Syntax:

```
set protocols mpls-rsvp log fsm egress
```

Syntax:

```
delete protocols mpls-rsvp log fsm egress
```

Syntax:

```
show protocols mpls-rsvp log fsm egress
```

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            fsm {
                egress
            }
        }
    }
}
```

Use this command to log egress state machine events.

Use the `set` form of this command to log egress state machine events.

Use the `delete` form of this command to delete the configuration to log egress state machine events.

Use the `show` form of this command to display the configuration to log egress state machine events.

protocols mpls-rsvp log fsm ingress

Logs ingress state machine events.

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
set protocols mpls-rsvp log fsm ingress
```

Syntax:

```
delete protocols mpls-rsvp log fsm ingress
```

Syntax:

```
show protocols mpls-rsvp log fsm ingress
```

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            fsm {
                ingress
            }
        }
    }
}
```

Use this command to log ingress state machine events.

Use the `set` form of this command to log ingress state machine events.

Use the `delete` form of this command to delete the configuration to log ingress state machine events.

Use the `show` form of this command to display the configuration to log ingress state machine events.

protocols mpls-rsvp log fsm transit downstream

Logs transit downstream state machine events.

Syntax:

```
set protocols mpls-rsvp log fsm transit downstream
```

Syntax:

```
delete protocols mpls-rsvp log fsm transit downstream
```

Syntax:

```
show protocols mpls-rsvp log fsm transit downstream
```

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            fsm {
                transit {
                    downstream
                }
            }
        }
    }
}
```

Use this command to log transit downstream state machine events.

Use the `set` form of this command to log transit downstream state machine events.

Use the `delete` form of this command to delete the configuration to log transit downstream state machine events.

Use the `show` form of this command to display the configuration to log transit downstream state machine events.

protocols mpls-rsvp log fsm transit upstream

Logs transit upstream state machine events.

Syntax:
```
set protocols mpls-rsvp log fsm transit upstream
```

Syntax:
```
delete protocols mpls-rsvp log fsm transit upstream
```

Syntax:
```
show protocols mpls-rsvp log fsm transit upstream
```

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            fsm {
                transit {
                    upstream
                }
            }
        }
    }
}
```

Use this command to log transit upstream state machine events.

Use the `set` form of this command to log transit upstream state machine events.

Use the `delete` form of this command to delete the configuration to log transit upstream state machine events.

Use the `show` form of this command to display the configuration to log transit upstream state machine events.

protocols mpls-rsvp log nsm

Logs interactions with the NSM.

Syntax:
```
set protocols mpls-rsvp log nsm
```

Syntax:
```
delete protocols mpls-rsvp log nsm
```

Syntax:
```
show protocols mpls-rsvp log nsm
```

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            nsm
        }
    }
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
}
```

Use this command to log interactions with the Network Services Module (NSM).

Use the `set` form of this command to log interactions with the NSM.

Use the `delete` form of this command to delete the configuration to log interactions with the NSM.

Use the `show` form of this command to display the configuration to log interactions with the NSM.

protocols mpls-rsvp log packet

Logs RSVP-TE packet transmission and reception.

Syntax:
`set protocols mpls-rsvp log packet`

Syntax:
`delete protocols mpls-rsvp log packet`

Syntax:
`show protocols mpls-rsvp log packet`

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            packet
        }
    }
}
```

Use this command to log RSVP-TE packet transmission and reception.

Use the `set` form of this command to log RSVP-TE packet transmission and reception.

Use the `delete` form of this command to delete the configuration to log RSVP-TE packet transmission and reception.

Use the `show` form of this command to display the configuration to log RSVP-TE packet transmission and reception.

protocols mpls-rsvp log rib

Logs interactions with the RIB.

Syntax:
`set protocols mpls-rsvp log rib`

Syntax:
`delete protocols mpls-rsvp log rib`

Syntax:
`show protocols mpls-rsvp log rib`

Configuration mode

```
protocols {
    mpls-rsvp {
        log {
            rib
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
        }
    }
}
```

Use this command to log interactions with the Routing Information Base (RIB).

Use the `set` form of this command to log interactions with the RIB.

Use the `delete` form of this command to delete the configuration to log interactions with the RIB.

Use the `show` form of this command to display the configuration to log interactions with the RIB.

protocols mpls-rsvp neighbors neighbor

Configures a neighbor to exchange hello packets.

Syntax:
`set protocols mpls-rsvp neighbors` neighbor *x.x.x.x*

Syntax:
`delete protocols mpls-rsvp neighbors` neighbor *x.x.x.x*

Syntax:
`show protocols mpls-rsvp neighbors` neighbor *x.x.x.x*

x.x.x.x
The address of the neighbor.

Configuration mode

```

protocols {
    mpls-rsvp {
        interfaces  {
            neighbors {
                neighbor <x.x.x.x>
            }
        }
    }
}
```

Use this command to configure a neighbor to exchange hello packets with so that the configuration can be used for link failure detection.

Use the `set` form of this command to configure a neighbor to exchange hello packets.

Use the `delete` form of this command to delete the configuration of a neighbor set to exchange hello packets.

Use the `show` form of this command to display the configuration of a neighbor set to exchange hello packets.

Example: Example

The following example shows how to enable RSVP-TE on the neighbor at 10.10.10.10.

```
vyatta@R1# set protocols mpls-rsvp neighbors neighbor 10.10.10.10
```

protocols mpls-rsvp tunnels tunnel

Configures the head-end for an RSVP-TE tunnel.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name*

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name*

name
The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name>
        }
    }
}
```

Use this command to configure the head-end for an RSVP-TE tunnel. No sessions will be established unless a destination is configured.

Use the `set` form of this command to configure the head-end for an RSVP-TE tunnel.

Use the `delete` form of this command to delete the configuration of the head-end for an RSVP-TE tunnel.

Use the `show` form of this command to display the configuration of the head-end for an RSVP-TE tunnel.

Example: Example

The following example shows how to configure the head-end for an RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2
```

protocols mpls-rsvp tunnels tunnel autoroute-announce

Enables calculation of routes over TE-tunnels by IGPs.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce

name
The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                autoroute-announce
            }
        }
    }
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
}
```

Use this command to enable calculation of routes over TE-tunnels by Interior Gateway Protocols (IGPs). By default, only the route to the tunnel destination is added to the RIB. When calculation of routes over TE-tunnels is enabled, IGPs treat a tunnel as a single link when calculating routes.

Use the `set` form of this command to enable calculation of routes over TE-tunnels by IGPs.

Use the `delete` form of this command to delete the configuration that enables calculation of routes over TE-tunnels by IGPs.

Use the `show` form of this command to display the configuration that enables calculation of routes over TE-tunnels by IGPs.

Example: Example

The following example shows how to allow IGPs to calculate routes over the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 autoroute-announce
```

protocols mpls-rsvp tunnels tunnel autoroute-announce absolute-metric

Configures a fixed metric to be used for the tunnel by the IGP for the shortest path calculations.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce absolute-metric *1..65535*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce absolute-metric *1..65535*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce absolute-metric *1..65535*

name

The tunnel name.

1..65535

The fixed metric value, which can be from 1 through 65535.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                autoroute-announce {
                    absolute-metric <1..65535>
                }
            }
        }
    }
}
```

Use this command to configure a fixed metric to be used for the specified tunnel by the Interior Gateway Protocols (IGPs) for its shortest path calculations. The default is the IGP metric.

Use the `set` form of this command to configure a fixed metric to be used for the tunnel by the IGP for its shortest path calculations.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use the `delete` form of this command to delete the configuration of a fixed metric to be used for the tunnel by the IGP for its shortest path calculations.

Use the `show` form of this command to display the configuration that sets a fixed metric to be used for the tunnel by the IGP for its shortest path calculations.

Example: Example

The following example shows how to configure 6000 as the IGP metric to be used by the IGP for the shortest path calculations for the RSVP-TE tunnel named pe1-pe2 tunnel.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 autoroute-announce absolute-metric
 6000
```

protocols mpls-rsvp tunnels tunnel autoroute-announce relative-metric

Configures a metric that is relative to the IGP metric.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce relative-metric *-65535..65535*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce relative-metric *-65535..65535*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* autoroute-announce relative-metric *-65535..65535*

name
The tunnel name.

-65535..65535
The value for a relative metric. The value can be from -65535 through 65535.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                autoroute-announce {
                    relative-metric <-65535..65535>
                }
            }
        }
    }
}
```

Use this command to configure a metric that is relative to the IGP metric for the specified tunnel.

The `relative-metric` command is mutually exclusive with the `absolute-metric` command.

Use the `set` form of this command to configure a metric that is relative to the IGP metric.

Use the `delete` form of this command to delete the configuration of a metric that is relative to the IGP metric.

Use the `show` form of this command to display the configuration that sets a metric that is relative to the IGP metric.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to configure 3000 as the metric that is relative to the IGP metric to be used for the RSVP-TE tunnel named pe1-pe2 tunnel by the IGP for its shortest path calculations.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 autoroute-announce relative-metric
 3000
```

protocols mpls-rsvp tunnels tunnel destination

Configures the address of the tail end of the tunnel.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* destination *x.x.x.x*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* destination *x.x.x.x*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* destination *x.x.x.x*

name
The tunnel name.

x.x.x.x
The address of the tail end of the tunnel.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                destination <x.x.x.x>
            }
        }
    }
}
```

Use this command to configure the address of the tail end of the specified tunnel. This is required in order to establish any sessions.

Use the `set` form of this command to configure the address of the tail end of the tunnel.

Use the `delete` form of this command to delete the configuration of the address of the tail end of the tunnel.

Use the `show` form of this command to display the configuration of the address of the tail end of the tunnel.

Example: Example

The following example shows how to configure 11.11.11.11 as the address of the tail end of the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 destination 11.11.11.11
```

protocols mpls-rsvp tunnels tunnel primary bandwidth

Configures the bandwidth to be reserved along the primary path.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary bandwidth *1-10000000000*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary bandwidth *1-10000000000*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary bandwidth *1-10000000000*

name

The tunnel name.

1-10000000000

The bandwidth reserved for the tunnel. The bandwidth can be from 1 through 10000000000.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    bandwidth <1-10000000000>
                }
            }
        }
    }
}
```

Use this command to configure the bandwidth to be reserved along the primary path. The interfaces must be configured with a maximum reservable bandwidth before tunnels can reserve bandwidth on the interfaces. By default, no bandwidth is reserved.

Use the `set` form of this command to configure the bandwidth to be reserved along the primary path.

Use the `delete` form of this command to delete the configuration of the bandwidth to be reserved along the primary path.

Use the `show` form of this command to display the configuration of the bandwidth to be reserved along the primary path.

Example: Example

The following example shows how to configure 1000000 as bandwidth to be reserved along the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary bandwidth 1000000
```

protocols mpls-rsvp tunnels tunnel primary explicit-path

Configures the primary path on the specified tunnel to follow the specified explicit path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary explicit-path *name*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary explicit-path *name*

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

`show protocols mpls-rsvp tunnels` tunnel *name* primary explicit-path *name*

name

The tunnel name.

name

The name of the explicit path.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    explicit-path <name>
                }
            }
        }
    }
}
```

Use this command to configure the primary path on the specified tunnel to follow the specified explicit path.

The primary path will follow the full or partial explicit path, which must be defined in the globals section. The default is to follow the CSPF calculated path.

Use the `set` form of this command to configure the primary path to follow the specified explicit path.

Use the `delete` form of this command to delete the configuration of the primary path to follow the specified explicit path.

Use the `show` form of this command to display the configuration of the primary path to follow the specified explicit path.

Example: Example

The following example shows how to configure the primary path for the RSVP-TE tunnel named pe1-pe2 to follow the explicit path named path1.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary explicit-path path1
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute

Enables fast-reroute on the primary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute
                }
            }
        }
    }
}
```

Use this command to enable fast-reroute for the primary path. This will establish node or link protection at each hop along the primary path.

Use the `set` form of this command to enable fast-reroute.

Use the `delete` form of this command to delete the configuration that enables fast-reroute.

Use the `show` form of this command to display the configuration that enables fast-reroute.

Example: Example

The following example shows how to configure the fast-reroute detour on the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute bandwidth

Enables fast-reroute with reserved bandwidth.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute bandwidth *1-10000000000*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute bandwidth *1-10000000000*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute bandwidth *1-10000000000*

name
The tunnel name.

1-10000000000
The bandwidth to be reserved along the detour. The bandwidth can be from 1 through 10000000000.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute {
                        bandwidth <1-10000000000>
                    }
                }
            }
        }
    }
}
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use this command to enable fast-reroute detours with reserved bandwidth. The interfaces must be configured with a maximum reservable bandwidth before tunnels can reserve bandwidth on the interfaces. By default, no bandwidth is reserved.

Use the `set` form of this command to enable fast-reroute with reserved bandwidth.

Use the `delete` form of this command to delete the configuration of the bandwidth to be reserved for the fast-reroute.

Use the `show` form of this command to display the configuration of the bandwidth to be reserved for the fast-reroute.

Example: Example

The following example shows how to configure 1000000 as bandwidth to be reserved for the fast-reroute detour for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute bandwidth
 1000000
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute node-protection-desired

Requests node protection along the path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute node-protection-desired

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute node-protection-desired

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute node-protection-desired

name
The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute {
                        node-protection-desired
                    }
                }
            }
        }
    }
}
```

Use this command to request node protection along the path for the fast-reroute detour. The default is link protection.

Use the `set` form of this command to request node protection along the path.

Use the `delete` form of this command to delete the configuration of the node protection along the path.

Use the `show` form of this command to display the configuration of the node protection along the path.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to configure node protection along the path for the fast-reroute detour for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute node-
protection-desired
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute path-selection affinities constraints include-any

Creates a list of affinities to include for the fast reroute detour.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints include-any

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints include-any

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints include-any

name
The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute {
                        path-selection {
                            affinities {
                                constraints include-any
                            }
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to include for the fast reroute detour for the primary path. The path will only be able to use interfaces bound to the included administrative groups

Use the `set` form of this command to create a list of affinities to include for the primary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to include for the primary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to include for the primary path.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to create a list of affinities to include for the fast reroute detour for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute path-selection
 affinities constraints include-any
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute path-selection affinities constraints include-any affinity-names

Creates a list of affinities to include for the specified administrative group for the fast reroute detour.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints include-any affinity-names *name*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints include-any affinity-names *name*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints include-any affinity-names *name*

name

The tunnel name.

name

The name of the administrative group.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute {
                        path-selection {
                            affinities {
                                constraints include-any {
                                    affinity-names <name>
                                }
                            }
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to include for the specified administrative group on the primary path fast reroute detour for the specified tunnel. The path will only be able to use interfaces bound to the included administrative groups.

Use the `set` form of this command to create a list of affinities to include for the specified administrative group on the primary path fast reroute detour.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to include for the specified administrative group on the primary path fast reroute detour.

Use the `show` form of this command to display the configuration that creates a list of affinities to include for the specified administrative group on the primary path fast reroute detour.

Example: Example

The following example shows how to create a list of affinities to include for the administrative group named group8 on the primary path fast reroute detour for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute path-selection
 affinities constraints include-any affinity-names group8
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute path-selection affinities constraints exclude-any

Creates a list of affinities to exclude for the fast reroute detour.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints exclude-any

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints exclude-any

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints exclude-any

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute {
                        path-selection {
                            affinities {
                                constraints exclude-any
                            }
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to exclude on the fast reroute detour for the primary path. The path will not be able to use interfaces bound to the excluded administrative groups.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use the `set` form of this command to create a list of affinities to exclude for the fast reroute detour.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to exclude for the fast reroute detour.

Use the `show` form of this command to display the configuration that creates a list of affinities to exclude for the fast reroute detour.

Example: Example

The following example shows how to create a list of affinities to exclude on the fast reroute detour for the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute path-selection
 affinities constraints exclude-any
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute path-selection affinities constraints exclude-any affinity-names

Creates a list of affinities to exclude for the specified administrative group on the fast reroute detour for the primary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints exclude-any affinity-names *name*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints exclude-any affinity-names *name*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute path-selection affinities constraints exclude-any affinity-names *name*

name

The tunnel name.

name

The name of the administrative group.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute {
                        path-selection {
                            affinities {
                                constraints exclude-any {
                                    affinity-names <name>
                                }
                            }
                        }
                    }
                }
            }
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
        }
    }
}
```

Use this command to create a list of affinities to exclude for the specified administrative group for the fast reroute detour on the primary path for the specified tunnel. The path will not be able to use the interfaces bound to the excluded administrative groups.

Use the `set` form of this command to create a list of affinities to exclude for the specified administrative group on the fast reroute detour.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to exclude for the specified administrative group on the fast reroute detour.

Use the `show` form of this command to display the configuration that creates a list of affinities to exclude for the specified administrative group on the fast reroute detour.

Example: Example

The following example shows how to create a list of affinities to exclude for the administrative group named group8 on the fast reroute detour for the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute path-selection
 affinities constraints exclude-any affinity-name group8
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute priority hold

Configures a hold priority for a session on a fast reroute detour.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute priority hold *0-7*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute priority hold *0-7*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute priority hold *0-7*

name

The tunnel name.

0-7

The hold priority for the session. The priority can be zero through seven, where zero is the highest priority.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    primary {
                        fast-reroute {
                            hold <0-7>
                        }
                    }
                }
            }
        }
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
    }
}
```

Use this command to configure a hold priority for this session on the fast reroute detour. The default is 0 (highest), which means it will not be pre-empted by higher priority sessions.

Use the `set` form of this command to configure a hold priority for this session on the fast reroute detour.

Use the `delete` form of this command to delete the configuration of a hold priority for this session on the fast reroute detour.

Use the `show` form of this command to display the configuration of a hold priority for this session on the fast reroute detour.

Example: Example

The following example shows how to set the hold priority to 7 (the lowest priority) for the session along the fast reroute detour for the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute priority hold
 7
```

protocols mpls-rsvp tunnels tunnel primary fast-reroute priority setup

Configures a setup priority for a session on a fast reroute detour.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute priority setup *0-7*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute priority setup *0-7*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute priority setup *0-7*

name

The tunnel name.

0-7

The setup priority for the session. The priority can be zero through seven, where zero is the highest priority.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute {
                        priority {
                            setup <0-7>
                        }
                    }
                }
            }
        }
    }
}
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use this command to configure a setup priority for this session on a fast reroute detour. The default is 7 (lowest), which means the session cannot pre-empt lower priority sessions.

Use the `set` form of this command to configure a setup priority for this session on a fast reroute detour.

Use the `delete` form of this command to delete the configuration of a setup priority for this session on a fast reroute detour.

Use the `show` form of this command to display the configuration of a setup priority for this session on a fast reroute detour.

Example: Example

The following example shows how to set the setup priority to zero for the session along the fast reroute detour for the primary path for the RSVP-TE tunnel named pe1-pe2. Zero is the highest priority, which means this session can pre-empt lower priority sessions.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute priority setup
 0
```

protocols mpls-rsvp tunnels tunne primary fast-reroute protection one-to-one

Configures one-to-one protection along the path for the fast-reroute detour.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute protection one-to-one

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute protection one-to-one

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary fast-reroute protection one-to-one

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    fast-reroute {
                        protection one-to-one
                    }
                }
            }
        }
    }
}
```

Use this command to configure one-to-one protection along the path for the fast-reroute detour. This is the default.

Use the `set` form of this command to request node protection along the path.

Use the `delete` form of this command to delete the configuration of the node protection along the path.

Use the `show` form of this command to display the configuration of the node protection along the path.

Example: Example

The following example shows how to configure one-to-one protection along the path for the fast-reroute detour for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary fast-reroute protection
 one-to-one
```

protocols mpls-rsvp tunnels tunnel primary lockdown

Locks the primary path on the specified tunnel.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary lockdown

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary lockdown

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary lockdown

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    lockdown
                }
            }
        }
    }
}
```

Use this command to lock the primary path on the specified tunnel. When the primary path is locked down, the primary path cannot be re-optimized.

Use the `set` form of this command to lock the primary path.

Use the `delete` form of this command to delete the configuration to lock the primary path.

Use the `show` form of this command to display the configuration to lock the primary path.

Example: Example

The following example shows how to lock the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary lockdown
```

protocols mpls-rsvp tunnels tunnel <name> primary no-cspf

Disables the use of CSPF for the primary path option.

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

`set protocols mpls-rsvp tunnels tunnel` *name* **`primary no-cspf`**

Syntax:
`delete protocols mpls-rsvp tunnels tunnel` *name* **`primary no-cspf`**

Syntax:
`show protocols mpls-rsvp tunnels tunnel` *name* **`primary no-cspf`**

`tunnel` ***name***
: The name of a TE tunnel.

Configuration mode

```
protocols {
        mpls-rsvp {
                   tunnels {
                              tunnel name {
                                       primary {
                                              no-cspf
                                       }
                              }
                   }
        }
}
```

Use the `set` form of this command to disable the use of CSPF for the primary path option.

Use the `delete` form of this command to enable the use of CSPF for the path option.

Use the `show` form of this command to display the CSPF configuration for the path option.

protocols mpls-rsvp tunnels tunnel primary path-selection affinities constraints include-any

Creates a list of affinities to include for the path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints include-any

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints include-any

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints include-any

name
: The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    path-selection {
                        affinities {
                            constraints include-any
                        }
                    }
                }
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
                }
            }
        }
    }
```

Use this command to create a list of affinities to include for the primary path. The path will only be able to use interfaces bound to the included administrative groups.

Use the `set` form of this command to create a list of affinities to include for the primary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to include for the primary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to include for the primary path.

Example: Example

The following example shows how to create a list of affinities to include for the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection affinities
 constraints include-any
```

protocols mpls-rsvp tunnels tunnel primary path-selection affinities constraints include-any affinity-names

Creates a list of affinities to include for the specified administrative group on the primary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints include-any affinity-names *name*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints include-any affinity-names *name*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints include-any affinity-names *name*

name

The tunnel name.

name

The name of the administrative group.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    path-selection {
                        affinities {
                            constraints include-any {
                                affinity-names <name>
                            }
                        }
                    }
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to include for the specified administrative group on the primary path for the specified tunnel. The path can only use interfaces bound to the included administrative groups.

Use the `set` form of this command to create a list of affinities to include for the specified administrative group on the primary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to include for the specified administrative group on the primary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to include for the specified administrative group on the primary path.

Example: Example

The following example shows how to create a list of affinities to include for the administrative group named group8 on the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection affinities
 constraints include-any affinity-names group8
```

protocols mpls-rsvp tunnels tunnel primary path-selection affinities constraints exclude-any

Creates a list of affinities to exclude for the path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints exclude-any

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints exclude-any

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints exclude-any

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    path-selection {
                        affinities {
                            constraints exclude-any
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to exclude for the primary path. The path will not be able to use interfaces bound to the excluded administrative groups.

Use the `set` form of this command to create a list of affinities to exclude for the primary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to exclude for the primary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to exclude for the primary path.

Example: Example

The following example shows how to create a list of affinities to exclude for the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection affinities
 constraints exclude-any
```

protocols mpls-rsvp tunnels tunnel primary path-selection affinities constraints exclude-any affinity-names

Creates a list of affinities to exclude for the specified administrative group on the primary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints exclude-any affinity-names *name*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints exclude-any affinity-names *name*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary path-selection affinities constraints exclude-any affinity-names *name*

name

The tunnel name.

name

The name of the administrative group.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    path-selection {
                        affinities {
                            constraints exclude-any {
                                affinity-names <name>
                            }
                        }
                    }
                }
            }
        }
    }
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
}
```

Use this command to create a list of affinities to exclude for the specified administrative group on the primary path for the specified tunnel. The path will not be able to use the interfaces bound to the excluded groups.

Use the `set` form of this command to create a list of affinities to exclude for the specified administrative group on the primary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to exclude for the specified administrative group on the primary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to exclude for the specified administrative group on the primary path.

Example: Example

The following example shows how to create a list of affinities to exclude for the administrative group named group8 on the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection affinities
 constraints exclude-any affinity-names group8
```

protocols mpls-rsvp tunnels tunnel primary path-selection hop-limit

Limits primary paths to those that do not exceed a number of hops for CSPF.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary path-selection hop-limit *value*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary path-selection hop-limit *value*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary path-selection hop-limit *value*

name

The tunnel name.

value

The maximum number of hops allowed.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    path-selection {
                        hop-limit <value>
                    }
                }
            }
        }
    }
}
```

Use this command to create a hop-limit constraint that limits CSPF to paths that do not exceed the set number of hops.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Use the `set` form of this command to create a hop-limit constraint for primary paths.

Use the `delete` form of this command to delete a hop-limit constraint for primary paths.

Use the `show` form of this command to display the configuration that creates a hop-limit constraint for primary paths.

Example: Example

The following example shows how to limit the primary path to 12 hops for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary path-selection hop-limit 12
```

protocols mpls-rsvp tunnels tunnel primary priority hold

Configures a hold priority for a session.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary priority hold *0-7*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary priority hold *0-7*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary priority hold *0-7*

name

The tunnel name.

0-7

The hold priority for the session. The priority can be zero through seven, where zero is the highest priority.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    priority {
                        hold <0-7>
                    }
                }
            }
        }
    }
}
```

Use this command to configure a hold priority for this session. The default is 0 (highest), which means it will not be pre-empted by higher priority sessions.

Use the `set` form of this command to configure a hold priority for this session.

Use the `delete` form of this command to delete the configuration of a hold priority for this session.

Use the `show` form of this command to display the configuration of a hold priority for this session.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to set the hold priority to 7 (the lowest priority) for the session along the primary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary priority hold 7
```

protocols mpls-rsvp tunnels tunnel primary priority setup

Configures a setup priority for a session.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary priority setup *0-7*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary priority setup *0-7*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary priority setup *0-7*

name

The tunnel name.

0-7

The setup priority for the session. The priority can be zero through seven, where zero is the highest priority.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    priority {
                        setup <0-7>
                    }
                }
            }
        }
    }
}
```

Use this command to configure a setup priority for this session. The default is 7 (lowest), which means the session cannot pre-empt lower priority sessions.

Use the `set` form of this command to configure a setup priority for this session.

Use the `delete` form of this command to delete the configuration of a setup priority for this session.

Use the `show` form of this command to display the configuration of a setup priority for this session.

Example: Example

The following example shows how to set the setup priority to zero for the session along the primary path for the RSVP-TE tunnel named pe1-pe2. Zero is the highest priority, which means this session can pre-empt lower priority sessions.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary priority setup 0
```

protocols mpls-rsvp tunnels tunnel primary record-route record-label

Records the labels used along the primary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* primary record-route record-label

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* primary record-route record-label

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* primary record-route record-label

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                primary {
                    record-route {
                        record-label
                    }
                }
            }
        }
    }
}
```

Use this command to record the labels in use along the primary path. The labels used along the path are displayed using the `show mpls rsvp session detail` command.

For example, here is a sample of the default `show mpls rsvp session detail output` output:

```
Record route: <var> 10.10.11.6 10.10.12.7
```

Compared with the `show mpls rsvp session detail output` output when this command is configured:

```
Record route: <var> 10.10.11.6 (52480) 10.10.12.7 (3)
```

Use the `set` form of this command to record the labels in use along the primary path.

Use the `delete` form of this command to remove the configuration to record the labels in use along the primary path.

Use the `show` form of this command to display the configuration to record the labels in use along the primary path.

Example: Example

The following example shows how to record the labels in use along the primary path for the RSVP-TE tunnel named pe1-pe2.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 primary record-route record-label
```

protocols mpls-rsvp tunnels tunnel secondary bandwidth

Configures the bandwidth to be reserved along the secondary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary bandwidth *1-10000000000*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary bandwidth *1-10000000000*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary bandwidth *1-10000000000*

name

The tunnel name.

1-10000000000

The bandwidth reserved for the tunnel. The bandwidth can be from 1 through 10000000000.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    bandwidth <1-10000000000>
                }
            }
        }
    }
}
```

Use this command to configure the bandwidth to be reserved along the secondary path. The interfaces must be configured with a maximum reservable bandwidth before tunnels can reserve bandwidth on the interfaces. By default, no bandwidth is reserved.

Use the `set` form of this command to configure the bandwidth to be reserved along the secondary path.

Use the `delete` form of this command to delete the configuration of the bandwidth to be reserved along the secondary path.

Use the `show` form of this command to display the configuration of the bandwidth to be reserved along the secondary path.

Example: Example

The following example shows how to configure 1000000 as the bandwidth to be reserved along the secondary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary bandwidth 1000000
```

protocols mpls-rsvp tunnels tunnel secondary explicit-path

Configures the secondary path on the specified tunnel to follow the specified explicit path.

Syntax:

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

`set protocols mpls-rsvp tunnels` tunnel *name* secondary explicit-path *name*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary explicit-path *name*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary explicit-path *name*

name

The tunnel name.

name

The name of the explicit path.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    explicit-path <name>
                }
            }
        }
    }
}
```

Use this command to configure the secondary path on the specified tunnel to follow the specified explicit path.

The secondary path will follow the full or partial explicit path, which must be defined in the globals section. The default is to follow the CSPF calculated path.

Use the `set` form of this command to configure the secondary path to follow the specified explicit path.

Use the `delete` form of this command to delete the configuration of the secondary path to follow the specified explicit path.

Use the `show` form of this command to display the configuration of the secondary path to follow the specified explicit path.

Example: Example

The following example shows how to configure the secondary path for the RSVP-TE tunnel named pe1-pe2 to follow the explicit path named path1.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary explicit-path path1
```

protocols mpls-rsvp tunnels tunnel secondary lockdown

Locks the secondary path on the specified tunnel.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary lockdown

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary lockdown

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary lockdown

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    lockdown
                }
            }
        }
    }
}
```

Use this command to lock the secondary path on the specified tunnel. When the secondary path is locked down, the secondary path cannot be re-optimized.

Use the `set` form of this command to lock the secondary path.

Use the `delete` form of this command to delete the configuration to lock the secondary path.

Use the `show` form of this command to display the configuration to lock the secondary path.

Example: Example

The following example shows how to lock the secondary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary lockdown
```

protocols mpls-rsvp tunnels tunnel <name> secondary no-cspf

Disables the use of CSPF for the secondary path option.

Syntax:

```
set protocols mpls-rsvp tunnels tunnel name secondary no-cspf
```

Syntax:

```
delete protocols mpls-rsvp tunnels tunnel name secondary no-cspf
```

Syntax:

```
show protocols mpls-rsvp tunnels tunnel name secondary no-cspf
```

`tunnel` ***name***

The name of a TE tunnel.

Configuration mode

```
protocols {
        mpls-rsvp {
                   tunnels {
                              tunnel name {
                                       secondary {
                                              no-cspf
                                     }
                              }
                   }
        }
}
```

Use the `set` form of this command to disable the use of CSPF for the secondary path option.

Use the `delete` form of this command to enable the use of CSPF for the secondary path option.

Use the `show` form of this command to display the CSPF configuration for the secondary path option.

protocols mpls-rsvp tunnels tunnel secondary path-selection affinities constraints include-any

Creates a list of affinities to include for the path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints include-any

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints include-any

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints include-any

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    path-selection {
                        affinities {
                            constraints include-any
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to include for the secondary path. The path will only be able to use interfaces bound to the included administrative groups

Use the `set` form of this command to create a list of affinities to include for the secondary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to include for the secondary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to include for the secondary path.

Example: Example

The following example shows how to create a list of affinities to include for the secondary path for the RSVP-TE tunnel named pe1-pe2.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary path-selection affinities
 constraints include-any
```

protocols mpls-rsvp tunnels tunnel secondary path-selection affinities constraints include-any affinity-names

Creates a list of affinities to include for the specified administrative group on the secondary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints include-any affinity-names *name*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints include-any affinity-names *name*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints include-any affinity-names *name*

name

The tunnel name.

name

The name of the administrative group.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    path-selection {
                        affinities {
                            constraints include-any {
                                affinity-names <name>
                            }
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to include for the specified administrative group on the secondary path for the specified tunnel. The path will only be able to use interfaces bound to the included administrative groups.

Use the `set` form of this command to create a list of affinities to include for the specified administrative group on the secondary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to include for the specified administrative group on the secondary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to include for the specified administrative group on the secondary path.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to create a list of affinities to include for the administrative group named group8 on the secondary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary path-selection affinities
 constraints include-any affinity-names group8
```

protocols mpls-rsvp tunnels tunnel secondary path-selection affinities constraints exclude-any

Creates a list of affinities to exclude for the path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints exclude-any

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints exclude-any

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints exclude-any

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    path-selection {
                        affinities {
                            constraints exclude-any
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to exclude for the secondary path. The path will not be able to use interfaces bound to the excluded administrative groups.

Use the `set` form of this command to create a list of affinities to exclude for the secondary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to exclude for the secondary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to exclude for the secondary path.

Example: Example

The following example shows how to create a list of affinities to exclude for the secondary path for the RSVP-TE tunnel named pe1-pe2.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary path-selection affinities
 constraints exclude-any
```

protocols mpls-rsvp tunnels tunnel secondary path-selection affinities constraints exclude-any affinity-names

Creates a list of affinities to exclude for the specified administrative group on the secondary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints exclude-any affinity-names *name*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints exclude-any affinity-names *name*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection affinities constraints exclude-any affinity-names *name*

name

The tunnel name.

name

The name of the administrative group.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    path-selection {
                        affinities {
                            constraints exclude-any {
                                affinity-names <name>
                            }
                        }
                    }
                }
            }
        }
    }
}
```

Use this command to create a list of affinities to exclude for the specified administrative group on the secondary path for the specified tunnel. The path will not be able to use the interfaces bound to the excluded administrative groups.

Use the `set` form of this command to create a list of affinities to exclude for the specified administrative group on the secondary path.

Use the `delete` form of this command to delete the configuration that creates a list of affinities to exclude for the specified administrative group on the secondary path.

Use the `show` form of this command to display the configuration that creates a list of affinities to exclude for the specified administrative group on the secondary path.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Example: Example

The following example shows how to create a list of affinities to exclude for the administrative group named group8 on the secondary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary path-selection affinities
 constraints exclude-any affinity-names group8
```

protocols mpls-rsvp tunnels tunnel secondary path-selection hop-limit

Limits secondary paths to those that do not exceed a number of hops for CSPF.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection hop-limit *value*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection hop-limit *value*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary path-selection hop-limit *value*

name

The tunnel name.

value

The maximum number of hops allowed.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    path-selection {
                        hop-limit <value>
                    }
                }
            }
        }
    }
}
```

Use this command to create a hop-limit constraint that limits CSPF to paths that do not exceed the set number of hops.

Use the `set` form of this command to create a hop-limit constraint for secondary paths.

Use the `delete` form of this command to delete a hop-limit constraint for secondary paths.

Use the `show` form of this command to display the configuration that creates a hop-limit constraint for secondary paths.

Example: Example

The following example shows how to limit the secondary path to 12 hops for the RSVP-TE tunnel named pe1-pe2.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary path-selection hop-limit
 12
```

protocols mpls-rsvp tunnels tunnel secondary priority hold

Configures a hold priority for a session.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary priority hold *0-7*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary priority hold *0-7*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary priority hold *0-7*

name

The tunnel name.

0-7

The hold priority for the session. The priority can be zero through seven, where zero is the highest priority.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    priority {
                        hold <0-7>
                    }
                }
            }
        }
    }
}
```

Use this command to configure a hold priority for this session. The default is 0 (highest), which means it will not be pre-empted by higher priority sessions.

Use the `set` form of this command to configure a hold priority for this session.

Use the `delete` form of this command to delete the configuration of a hold priority for this session.

Use the `show` form of this command to display the configuration of a hold priority for this session.

Example: Example

The following example shows how to set 7 (the lowest priority) as the hold priority for the session along the secondary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary priority hold 7
```

protocols mpls-rsvp tunnels tunnel secondary priority setup

Configures a setup priority for a session.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary priority setup *0-7*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary priority setup *0-7*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary priority setup *0-7*

name

The tunnel name.

0-7

The setup priority for the session. The priority can be zero through seven, where zero is the highest priority.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    priority {
                        setup <0-7>
                    }
                }
            }
        }
    }
}
```

Use this command to configure a setup priority for this session. The default is 7 (lowest), which means the session cannot pre-empt lower priority sessions.

Use the `set` form of this command to configure a setup priority for this session.

Use the `delete` form of this command to delete the configuration of a setup priority for this session.

Use the `show` form of this command to display the configuration of a setup priority for this session.

Example: Example

The following example shows how to set the setup priority to zero for the session along the secondary path for the RSVP-TE tunnel named pe1-pe2. Zero is the highest priority, which means this session can pre-empt lower priority sessions.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary priority setup 0
```

protocols mpls-rsvp tunnels tunnel secondary record-route record-label

Records the labels used along the secondary path.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* secondary record-route record-label

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* secondary record-route record-label

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* secondary record-route record-label

name

The tunnel name.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                secondary {
                    record-route {
                        record-label
                    }
                }
            }
        }
    }
}
```

Use this command to record the labels in use along the secondary path. The labels used along the path are displayed using the `show mpls rsvp session detail` command.

For example, here is a sample of the default `show mpls rsvp session detail output` output:

```
Record route: <var> 10.10.11.6 10.10.12.7
```

Compared with the `show mpls rsvp session detail output` output when this command is configured:

```
Record route: <var> 10.10.11.6 (52480) 10.10.12.7 (3)
```

Use the `set` form of this command to record the labels in use along the secondary path.

Use the `delete` form of this command to remove the configuration to record the labels in use along the secondary path.

Use the `show` form of this command to display the configuration to record the labels in use along the secondary path.

Example: Example

The following example shows how to record the labels in use along the secondary path for the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 secondary record-route record-label
```

protocols mpls-rsvp tunnels tunnel source

Configures the source address of the MPLS RSVP tunnel.

Syntax:
`set protocols mpls-rsvp tunnels` tunnel *name* source *x.x.x.x*

Syntax:
`delete protocols mpls-rsvp tunnels` tunnel *name* source *x.x.x.x*

Syntax:
`show protocols mpls-rsvp tunnels` tunnel *name* source *x.x.x.x*

name

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

The tunnel name.

x.x.x.x

The source address of the tunnel.

Configuration mode

```
protocols {
    mpls-rsvp {
        tunnels  {
            tunnel <name> {
                source <x.x.x.x>
            }
        }
    }
}
```

Use this command to configure the source address of the specified tunnel. This defaults to the address on the first interface.

Use the `set` form of this command to configure the source address of the tunnel.

Use the `delete` form of this command to delete the configuration of the source address of the tunnel.

Use the `show` form of this command to display the configuration of the source address of the tunnel.

Example: Example

The following example shows how to configure 11.11.11.10 as the source address of the RSVP-TE tunnel named pe1-pe2.

```
vyatta@R1# set protocols mpls-rsvp tunnels tunnel pe1-pe2 source 11.11.11.10
```

reset mpls rsvp tunnel all [primary | secondary]

Resets sessions for all tunnels.

Syntax:
`reset mpls rsvp tunnel all` [primary | secondary]

Operational mode

Use this command to reset sessions for all tunnels. The default is to tear down and re-establish all sessions. Optionally, you can reset only the primary sessions or only the secondary sessions.

reset mpls rsvp tunnel ingress all | name

Resets ingress sessions for either all tunnels or a specific tunnel.

Syntax:
`reset mpls rsvp tunnel ingress` all | name *tunnel-name*

tunnel-name

The name of the specific tunnel.

Operational mode

Use this command to reset ingress sessions for either all tunnels or a specific tunnel.

reset mpls rsvp tunnel name <name> [primary | secondary]

Resets sessions for a specific tunnel.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`reset mpls rsvp tunnel name` *name* [primary | secondary]

tunnel-name
The name of the specific tunnel.

Operational mode

Use this command to reset sessions for a specific tunnel. The default is to tear down and re-establish all sessions. Optionally, you can reset only the primary sessions or only the secondary sessions.

reset mpls rsvp tunnel non-ingress all | name

Resets non-ingress sessions for either all tunnels or a specific tunnel.

Syntax:
`reset mpls rsvp tunnel non-ingress` all | name *tunnel-name*

tunnel-name
The name of the specific tunnel.

Operational mode

Use this command to reset non-ingress sessions (egress and transit sessions) for either all tunnels or a specific tunnel.

reset mpls rsvp tunnel reoptimize all | name [primary | secondary]

Resets optimization in the sessions for either all tunnels or a specific tunnel.

Syntax:
`reset mpls rsvp tunnel reoptimize` all | name *tunnel-name* [primary | secondary]

tunnel-name
The name of the specific tunnel.

Operational mode

Use this command to reoptimize the sessions for either all tunnels or a specific tunnel. The default is to reoptimize all sessions. Optionally, you can reset optimization in only the primary sessions or only the secondary sessions.

Reoptimization involves establishing a new make-before-break session if there is a better path available. No new session is established if there is no better path.

show mpls rsvp

Displays the global MPLS RSVP information.

Syntax:
`show mpls rsvp`

Operational mode

Use this command to display the global MPLS RSVP information.

The following example shows how to display global MPLS RSVP information.

```
vyatta@vyatta:~$ show mpls rsvp
RSVP Version                  : 1
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
Process uptime                : 5 minutes
Stagger timer                 : Not running
RSVP Refresh Reduction        : Enabled
RSVP Message Acknowledgement  : Disabled
Bundle Send                   : Disabled
NSM Connection                : Up
CSPF Connection IPv4          : Up
CSPF Connection IPv6          : Down
CSPF usage                    : Enabled
Reoptimization                : Disabled
RSVP Refresh Timer            : 30
Keep Multiplier               : 3
Acknowledgement Await Timeout : 10
Explicit-Null For Direct Conn : Disabled
Local Protection              : Disabled
Hello Receipt                 : Disabled
Hello Interval                : 2000
Hello Timeout                 : 7000
Loop detection                : Enabled (all interface)
Ingress                       : 5.5.5.5
Ingress                       : N/A (not in use)
Penultimate Hop Popping       : Enabled
Refresh PATH msg parsing      : Enabled
Refresh RESV msg parsing      : Enabled
Detour identification         : Sender-Template
Notification                  : Disabled
```

show mpls rsvp admin-groups

Displays configured administrative groups.

Syntax:

`show mpls rsvp admin-groups`

Operational mode

Use this command to display all the configured MPLS RSVP administrative groups.

The following example shows how to display the MPLS RSVP administrative groups (named red and blue in this example) and their respective associated values.

```
vyatta@vyatta:~$ show mpls rsvp admin-groups
Admin group detail:
 Value of 1 associated with admin group red
 Value of 2 associated with admin group blue
```

show mpls rsvp graceful-restart

Displays information about graceful-restart hello packets.

Syntax:

`show mpls rsvp graceful-restart` [*ip-address-of-neighbor*]

ip-address-of-neighbor

The IP address of an MPLS RSVP neighbor.

Operational mode

Use this command to display information about graceful-restart hello packets.

The following example shows the graceful restart information for MPLS RSVP.

```
 vyatta@vyatta:~$ show mpls rsvp graceful-restart
Graceful Restart: Enabled
Advertised Restart Time: 0 msec
Advertised (default) Recovery Time: 0 msec

  Remote addr: 192.166.1.2           Local addr: 192.166.1.1
  Interface: dp0p1s3
  Nbr State: Normal, Type: Reroute
  Nbr Hello State: Up
  LSPs being protected: 2
  Number of Hello messages sent: 84597
  Number of Hello messages received: 81036
  Last Hello message sent
    Elapsed time: 1 msec, Type: Ack
    Src instance: 0x45d18bc4, Dst instance: 0x5ade2a5b
    Restart time: 0 msec, Recovery time: 0 msec
  Last Hello message received
    Elapsed time: 1 msec, Type: Ack
    Src instance: 0x5ade2a5b, Dst instance: 0x45d18bc4
    Restart time: 0 msec, Recovery time: 0 msec

  Remote addr: 192.166.2.2           Local addr: 192.166.2.1
  Interface: dp0p1s1
  Nbr State: Normal, Type: Reroute
  Nbr Hello State: Up
  LSPs being protected: 3
  Number of Hello messages sent: 81519
  Number of Hello messages received: 74406
  Last Hello message sent
    Elapsed time: 6 msec, Type: Ack
    Src instance: 0x2, Dst instance: 0x7
    Restart time: 0 msec, Recovery time: 0 msec
  Last Hello message received
    Elapsed time: 3 msec, Type: Ack
    Src instance: 0x7, Dst instance: 0x2
    Restart time: 0 msec, Recovery time: 0 msec
```

The following example shows graceful restart information for an MPLS RSVP neighbor that does not have the capability for graceful restart.

```
 vyatta@vyatta:~$ show mpls rsvp graceful-restart 192.166.2.2
Graceful Restart: Enabled
Advertised Restart Time: 0 msec
Advertised (default) Recovery Time: 0 msec

  Remote addr: 192.166.2.2           Local addr: 192.166.2.1
  Interface: dp0p1s1
  Nbr State: Normal, Type: Reroute
  Nbr Hello State: Up
  LSPs being protected: 2
  Number of Hello messages sent: 92600
  Number of Hello messages received: 85487
  Last Hello message sent
    Elapsed time: 6 msec, Type: Ack
    Src instance: 0x2, Dst instance: 0x7
    Restart time: 0 msec, Recovery time: 0 msec
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
  Last Hello message received
    Elapsed time: 5 msec, Type: Ack
    Src instance: 0x7, Dst instance: 0x2
    Restart time: 0 msec, Recovery time: 0 msec
```

The following example shows graceful restart information for a neighbor that is capable of graceful restart.

```
 vyatta@vyatta:~$ show mpls rsvp graceful-restart 10.10.12.66
Graceful Restart: Enabled
Advertised Restart Time: 0 msec
Advertised (default) Recovery Time: 0 msec

  Remote addr: 10.10.12.66           Local addr: 10.10.12.6
  Interface: dp0p1s12
  Nbr State: Normal, Type: Graceful Restart
  Nbr Hello State: Up
  LSPs being protected: 1
  Number of Hello messages sent: 118
  Number of Hello messages received: 118
  Last Hello message sent
    Elapsed time: 6184 msec, Type: Ack
    Src instance: 0x3c, Dst instance: 0xd72f6b0c
    Restart time: 0 msec, Recovery time: 0 msec
  Last Hello message received
    Elapsed time: 6184 msec, Type: Req
    Src instance: 0xd72f6b0c, Dst instance: 0x3c
    Restart time: 60000 msec, Recovery time: 0 msec
```

show mpls rsvp interface

Displays which interfaces are MPLS RSVP-enabled or the RSVP parameters configured on a specified interface.

Syntax:

`show mpls rsvp interface` [*name*]

name

The name of the specified interface.

Operational mode

Use this command to display which interfaces are MPLS RSVP- enabled or to display the RSVP parameters configured on a specified interface.

The following example shows how to use the `show mpls rsvp interface` command to display which interfaces are MPLS RSVP-enabled.

```
 vyatta@vyatta:~$ show mpls rsvp interface
Interface     RSVP status   Interface Type
lo            Disabled      N/A
eth0          Disabled      N/A
dp0p1s6       Disabled      N/A
dp0p1s10      Enabled       Ethernet
dp0p1s11      Enabled       Ethernet
dp0p1s15      Enabled       Ethernet
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

The following example shows how to display which RSVP parameters are configured on a specified MPLS RSVP-enabled interface.

```
vyatta@vyatta:~$ show mpls rsvp interface dp0p1s11
Status                        : Enabled
Interface Index               : 10
Refresh Reduction usage       : Enabled
Message Acknowledgement       : Disabled
Bundle Buffer size            : 65532
Current Epoch Value           : 247115164
Primary IPv4 address          : 10.10.11.5
Primary IPv6 address          : fe80::5054:ff:fe00:511
Interface Type                : Ethernet
Administrative Group          : blue
Configured refresh time       : 30
Configured keep multiplier    : 3
Acknowledgement Await Timeout : 10
Hello Receipt                 : Disabled
Hello Interval                : 2000
Hello Timeout                 : 7000
Non IANA Hello exchange       : Disabled
```

show mpls rsvp neighbor

Displays a summary of all RSVP neighbors, or the list of active LSPs for a specific neighbor.

Syntax:

`show mpls rsvp neighbor` [*ip-address*]

ip-address

The address of a specific neighbor.

Operational mode

Use this command to display a summary of all MPLS RSVP neighbors, or the list of active LSPs for a specific neighbor.

The following example shows how to display a summary of all RSVP neighbors.

```
 vyatta@vyatta:~$ show mpls rsvp neighbor
IP Address          UpStrm LSP DnStrm LSP RefreshReduc Srefresh    In  Type      GraceRestart
192.166.2.1         1          0          Enabled      22s             Explicit  Incapable
192.166.3.1         0          0          Enabled      not running     Explicit  Incapable
192.166.4.2         0          0          Enabled      not running     Explicit  Incapable
192.166.8.2         0          1          Enabled      2s              Explicit  Incapable
```

The following example shows how to display the list of active LSPs for a specific neighbor.

```
vyatta@vyatta:~$ show mpls rsvp neighbor 192.166.2.1
 Nbr Hello State: Up
 Upstream LSPs: 0, Downstream LSPs: 0
 Neighbor supports Refresh Reduction, next SRefresh transmission in: 15s
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
Neighbor does not support Graceful Restart.
```

The following example shows how to display the list of active LSPs for a neighbor that is capable of grateful restart.

```
vyatta@vyatta:~$ show mpls rsvp neighbor 10.10.12.66
 Nbr Hello State: Up
 Upstream LSPs: 1, Downstream LSPs: 0
 Neighbor does not support Refresh Reduction, no SRefresh transmission scheduled
 Neighbor supports Graceful Restart, Restart time: 60000 msec, Recovery time: 0 msec

  Tunnel ID  LSP ID     Ingress             Egress              Type
  4676       2          66.66.66.66         2.2.2.2             Upstream
```

show mpls rsvp path

Displays the explicit path configuration.

Syntax:

`show mpls rsvp path` [*name*]

name

The name of the specified path.

Operational mode

Use this command to display the explicit path configuration for all paths or for a specified path.

The following example shows how to display the explicit path configuration for all paths.

```
 vyatta@vyatta:~$ show mpls rsvp path
Path name: p678, id: 2, hop-count: 2 type: mpls
 6.6.6.6 loose
 7.7.7.7 loose

Path name: p238, id: 1, hop-count: 3 type: mpls
 2.2.2.2 loose
 3.3.3.3 loose
 8.8.8.8 loose
```

The following example shows how to display the explicit path configuration for the path named "p678".

```
 vyatta@vyatta:~$ show mpls rsvp path p678
Path name: p678, id: 2, hop-count: 2 type: mpls
 6.6.6.6 loose
 7.7.7.7 loose
```

show mpls rsvp session

Displays a list of the MPLS RSVP sessions.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`show mpls rsvp session`

Operational mode

Use this command to display a list of the current MPLS RSVP sessions.

The following example shows how to display the list of current MPLS RSVP sessions to and from 8.8.8.8.

```
vyatta@vyatta:~$ show mpls rsvp session
Ingress RSVP:
To              From            State Pri Rt Style Labelin Labelout LSPname
8.8.8.8         5.5.5.5         Up    Yes  1  1 SE       -   53120 t8
8.8.8.8         10.10.10.5      Up    No   1  1 SE       -   52480 t8
Total 2 displayed, Up 2, Down 0.

Egress RSVP:
To              From            State Pri Rt Style Labelin Labelout LSPname
5.5.5.5         8.8.8.8         Up    Yes  1  1 SE       3        - t5
Total 1 displayed, Up 1, Down 0.
```

show mpls rsvp session count

Displays the total number of configured, ingress, egress, and transit sessions.

Syntax:
`show mpls rsvp session count`

Operational mode

Use this command to display the total number of:

- Sessions
- Sessions that are up
- Sessions that are down
- Ingress sessions
- Ingress sessions that are up
- Ingress sessions that are down
- Egress sessions
- Egress sessions that are up
- Egress sessions that are down
- Transit sessions
- Transit sessions that are up
- Transit sessions that are down

The following example shows how to display the number of current sessions.

```
vyatta@vyatta:~$ show mpls rsvp session count
Total configured: 3, Up 3, Down 0

Total ingress sessions: 2, Up 2, Down 0
Total transit sessions: 0, Up 0, Down 0
Total egress sessions: 1, Up 1, Down 0
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

show mpls rsvp session [egress | ingress | transit] [down | up] detail

Displays detailed session information.

Syntax:

`show mpls rsvp session` [egress | ingress | transit] [down | up] detail

Operational mode

Use this command to display detailed session information. The default is to show all sessions. You can limit the output to just egress, ingress, or transit sessions and/or sessions that are up or down.

The following example shows how to use the `show mpls rsvp session egress detail` command to display detailed information on the current egress session.

```
vyatta@vyatta:~$ show mpls rsvp session egress detail
Egress (Primary)
5.5.5.5
  From: 8.8.8.8, LSPstate: Up, LSPname: t5
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in:        3,  Label out: -
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5001, LSP Id: 101, Ext-Tunnel Id: 8.8.8.8
  Upstream: 10.10.15.2, dp0p1s15
  Path lifetime: 157 seconds (due in 120 seconds)
  Resv refresh: 30 seconds (due in 32975 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Hop Limit: 255
  Admin Groups:   Received Explicit Route Detail :
   10.10.15.5/32 strict
  Record route: 10.10.18.8 10.10.3.3 10.10.15.2 <self>
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
```

The output in the following example indicates that the use of CSPF is disabled.

```
vyatta@vyatta:~$ show mpls rsvp session detail
Ingress (Primary)
6.6.6.6
  From: 5.5.5.5, LSPstate: Up, LSPname: t6
  Ingress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  CSPF usage: Disabled
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

```
...
```

show mpls rsvp session name <name> [primary | secondary]

Displays detailed session information for a specific tunnel.

Syntax:

`show mpls rsvp session name` *name* [primary | secondary]

name

The name of the specific tunnel.

Operational mode

Use this command to display detailed session information for a specific tunnel. The default is to show all sessions. You can limit the output to just the primary or secondary (if configured) session.

The following example shows how to display detailed session information from the primary session in the tunnel named t8.

```
vyatta@vyatta:~$ show mpls rsvp session name t8 primary
Ingress (Primary)
8.8.8.8
  From: 5.5.5.5, LSPstate: Up, LSPname: t8
  Ingress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  Reoptimization: Disabled
  IGP-Shortcut: Disabled, LSP metric: 25
  LSP Protection: one-to-one
  Label in: -,  Label out:    53120
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5001, LSP Id: 101, Ext-Tunnel Id: 5.5.5.5
  Downstream: 10.10.15.2, dp0p1s15
  Path refresh: 30 seconds (RR enabled) (due in 28 seconds)
  Resv lifetime: 157 seconds (due in 138 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: p238 (in use)
  Configured Explicit Route Detail :
   2.2.2.2/32 loose
   3.3.3.3/32 loose
   8.8.8.8/32 loose
  Session Explicit Route Detail :
   10.10.15.2/32 strict
   10.10.3.3/32 strict
   10.10.18.8/32 strict
  Record route: <self> 10.10.15.2 10.10.3.3 10.10.18.8
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

show mpls rsvp statistics

Displays MPLS RSVP packet sent and received statistics.

Syntax:

```
show mpls rsvp statistics
```

Operational mode

Use this command to display the number of sent and received packets per MPLS RSVP packet type.

The following example shows how to display the sent and received statistics per MPLS RSVP packet type.

```
vyatta@vyatta:~$ show mpls rsvp statistics
  PacketType               Total
                    Sent      Received
  Path                 9            2
  PathErr              0            0
  PathTear             2            0
  Resv FF              0            0
  Resv WF              0            0
  Resv SE              2           10
  Resv Err             0            0
  ResvTear             0            0
  ResvConf             0            0
  Hello                0            0
  Bundle               0            0
  Ack                  0            0
  SRefresh           129          131
  Notify               0            0
```

show mpls rsvp summary-refresh

Displays sessions using summary refresh.

Syntax:

```
show mpls rsvp summary-refresh
```

Operational mode

Use this command to display the sessions that are using summary refresh.

The following example shows how to display the sessions using summary refresh.

```
 vyatta@vyatta:~$ show mpls rsvp summary-refresh
Neighbor Addr       Tunnel ID  LSP ID     Ingress             Egress
10.10.11.6          5001       101        10.10.10.5          8.8.8.8
10.10.15.2          5001       101        8.8.8.8             5.5.5.5
10.10.15.2          5001       101        5.5.5.5             8.8.8.8
```

show mpls rsvp tunnel

Displays a summary list of the configured MPLS RSVP tunnels.

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.

Syntax:
`show mpls rsvp tunnel`

Operational mode

Use this command to display a summary list of the MPLS RSVP tunnels.

The following example shows how to display a summary list of the configured MPLS RSVP tunnels.

```
 vyatta@vyatta:~$ show mpls rsvp tunnel
Trunk Name      Trunk ID  Type  # Sess  Egress Address(es)
N/A             5001      P2P   1       5.5.5.5
t8              5001      P2P   2       8.8.8.8
Total trunks configured: 2.
```

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property.