

# Point-to-Point LSPs

---

## In This Chapter

This section provides information about point-to-point LSPs (static, LDP and RSVP-TE).

Topics in this section include:

- [Applicability on page 1156](#)
- [Overview on page 1157](#)
- [Configuration on page 1161](#)
- [Conclusion on page 1182](#)

## Applicability

This configuration note is applicable to all of the 7750, 7450, and 7710 SR and ESS series. It was tested on release 8.0R4. There are no pre-requisites or conditions on the hardware for this configuration.

## Overview

Due to the connectionless nature of the network layer protocol IP, packets travel through the network on a hop-by-hop basis with routing decisions made at each node. As a result, hyper aggregation of data on certain links may occur and it may impact the provider's ability to provide guaranteed service levels across the network end-to-end. To address these shortcomings, MPLS (Multiprotocol Label Switching) was developed. The technology provides the capability to establish connection oriented paths, called Label Switched Paths (LSPs), over a connectionless (IP) network. The LSP offers a mechanism to engineer network traffic independently from the underlying network routing protocol (mostly IP) to improve the network resiliency and recovery options and to permit delivery of new services that are not readily supported by conventional IP routing techniques (Layer 2 IP VPNs). These benefits are essential for today's communication network explaining the wide deployment base of the MPLS technology.

RFC 3031, *Multiprotocol Label Switching Architecture*, specifies the MPLS architecture while this document describes the configuration and troubleshooting of point-to-point LSPs on Alcatel-Lucent SR and ESS series routers.

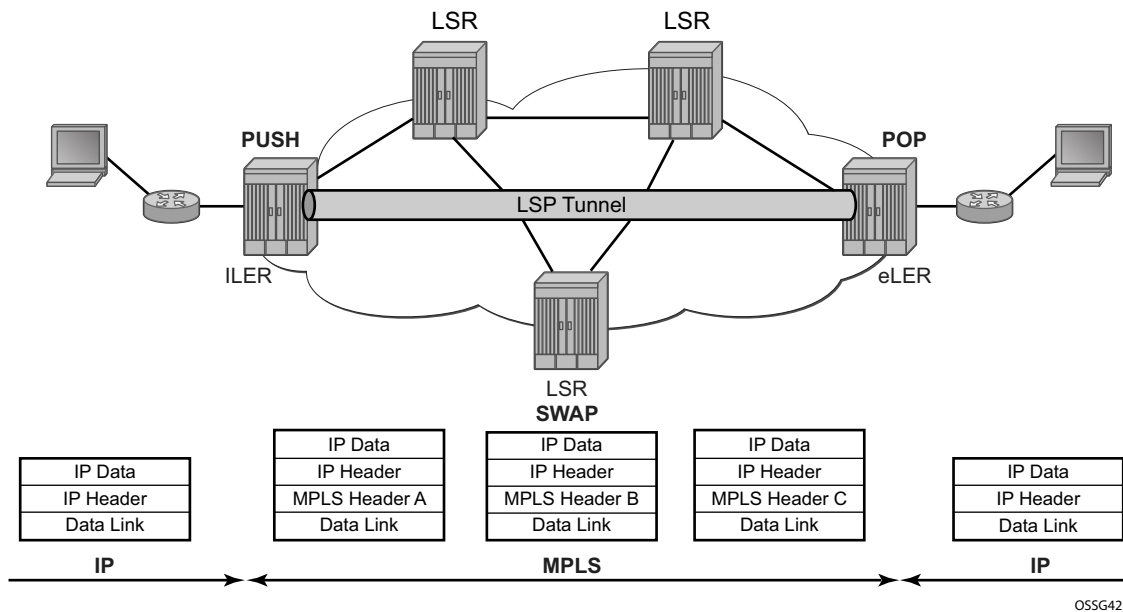
---

## Packet Forwarding

As a packet of a connectionless network layer protocol travels from one router to the next, each router in the network makes an independent forwarding decision by performing the following basic tasks: first analyzing the packet's header, then referencing the local routing table to find the longest match based on the destination address in the IP header, and finally sending out the packet on the selected interface. In other terms, the first function partitions the entire set of possible packets into a set of Forwarding Equivalence Classes (FECs). All packets associated to a particular FEC will be forwarded along the same logical path to the same destination. The second function maps each FEC to a next hop destination router. Each router along the packet's path performs these actions.

On the other hand, in MPLS the assignment of a particular packet to a particular FEC is done just once, as the packet enters the network. In turn the FEC is mapped to an LSP, which is pre-sigaled prior to any data flowing. An MPLS label, representing the FEC to which the packet is assigned, is attached to the packet (push operation) and once labeled the packet is forwarded to the next hop router along that LSP path. At subsequent hops, there is no further analysis of the packet's network layer header. Instead, the label is used as an index into a table which specifies the next hop and a new label. The old label is replaced with the new label (swap operation), and the packet is forwarded to its next hop. At the MPLS network egress, the label is removed from the packet (POP operation). If this router is the final destination (based on the remaining packet), the packet is handed to the receiving application (such as a VPLS domain). If this router is not the final destination of the packet, the packet will be sent into a new MPLS tunnel or forwarded by conventional IP forwarding towards the Layer 3 destination.

## Terminology



**Figure 166: Generic MPLS Network, MPLS Label Operations**

Figure 166 depicts a general network topology clarifying the MPLS-related terms. A Label Edge Router (LER) is a device at the edge of an MPLS network, with at least one interface outside the MPLS domain. A router is usually defined as an LER based on its position relative to a particular LSP. The MPLS router at the head-end of an LSP is called the ingress Label Edge Router (iLER). The MPLS router at the tail-end of an LSP is called the egress Label Edge Router (eLER). The iLER receives unlabeled packets from outside the MPLS domain then applies MPLS labels to the packets and forwards the labeled packets into the MPLS domain. The eLER receives labeled packets from the MPLS domain then removes the labels and forwards unlabeled packets outside the MPLS domain. Note that Release 8.0 added the ability for the eLER to signal an implicit-null label (numeric value 3). This informs the previous hop to send MPLS packets without an outer label and so is known as penultimate hop popping (PHP). This is also available when using static LSPs.

A Label Switching Router (LSR) is a device internal to an MPLS network, with all interfaces inside the MPLS domain. These devices switch labeled packets inside the MPLS domain. In the core of the network, LSR ignore the packet's network layer (IP) header and simply forward the packet using the MPLS label swapping mechanism.

A single LSP is uni-directional. In common practice, because the bi-directional nature of most traffic flows is implied, the term LSP often is used to define the pair of LSPs that enable the bidirectional flow. For ease of terminology and discussion however, the LSP in this document is referred to as a single entity.

## LSP Establishment

Prior to packet forwarding, the LSP must be established. In order to do so, labels need to be distributed for the path. Labels are usually distributed by a downstream router in the upstream direction (relative to the data flow). There are a number of ways used for label distribution.

- The label distribution can be done manually by the network administrator by configuring static LSPs. Although a high control level of the labels in use is achieved, the LSP cannot enjoy the resilience and recovery functionality the dynamic label signaling protocols can offer.
- LDP (RFC 5036, *LDP Specification*) can be considered as an extension to the network IGP. As routers become aware of new destination networks, they advertise labels in the upstream direction that will allow upstream routers to reach the destination.
- RSVP-TE (RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*) can also be used to signal LSPs across the network. RSVP-TE is used for traffic engineering when the ingress router wishes to create an LSP with specific constraints beyond the best route chosen by the IGP. RSVP-TE identifies the specific path desired for the LSP and may include resource requirements for the path.

The most important benefit of the label swapping mechanism RSVP-TE is its ability to map any type of user traffic to a LSP that has been specifically engineered to satisfy user traffic requirements. Customized LSPs may be created based on hop count or bandwidth requirements. They can even be routed through specific network links or nodes, as specified by the ingress node. This offers service providers precise control over the flow of traffic in their networks and results in a network that operates more efficiently and provides more predictable and scalable services.

## Testbed Topology

The network topology is displayed in [Figure 167](#). The setup consists of six 7750 nodes located in a single autonomous system.




Figure 167: MPLS Testbed Topology

## Configuration

As a general pre-requisite for the configuration of MPLS LSPs, a correctly working IGP is required<sup>1</sup>; either OSPF or IS-IS can be used for this purpose.

For LSPs that are set up manually or using RSVP-TE, the first step is to enable MPLS on:

- All network interfaces that will be used to carry LSPs
- System IP address

For manually configured LSPs, any interface used by the static LSP must be added into the MPLS protocol instance, even though RSVP is not actually used to signal labels. For router PE-1 this results in the following configuration:


```
A:PE-1# configure router mpls
A:PE-1>config>router>mpls# info
-----
        interface "system"
        exit
        interface "int-PE-1-PE-2"
        exit
        interface "int-PE-1-PE-4"
        exit
        no shutdown
-----
A:PE-1>config>router>mpls#
```

---

1. Static and strict path LSPs do not need an IGP.

## Manually Configured LSPs

As an example, a static LSP will be created starting from PE-1, running over PE-2 and PE-5, then terminating on PE-6 as depicted in Figure 168.



**Figure 168: Static LSP Running Over PE-1 PE-2 PE-5 PE-6**

For each node verify the available labels prior to beginning the configuration and verify the acceptable label range for use with static configurations.

```
A:PE-1# show router mpls label-range
=====
Label Ranges
=====
Label Type      Start Label    End Label      Aging          Total Available
-----
Static-lsp      32             1023           -              992
Static-svc      2048           18431          -              16384
Dynamic         32768          131071         0              98304
=====
A:PE-1#
```



The label range for static LSPs extends from the value 32 to 1023. To ensure the labels have not already been allocated to another configuration, use the command:

```
A:PE-1# show router mpls label 32 1023 in-use
=====
MPLS Labels from 32 to 1023 (In-use)
=====
Label                Label Type          Label Owner
-----
In-use labels (Owner: All) in specified range : 0
In-use labels in entire range                 : 0
=====
A:PE-1#
```

As no LSPs are currently originating, passing through, or terminating in the node, no labels are in use and any label from the range 32 to 1023 is available for the static LSP. For the originating router PE-1, the label 100 will be used for the push operation on the interface towards PE-2.

Static LSPs are configured within the MPLS configuration context, but do not rely on dynamic label signaling.

The configuration of the MPLS static LSP head-end PE-1 contains:

- The system IP address of the destination router PE-6 (to).
- A push operation of the label 100.
- The interface address facing the current node of the next hop along the static path, which is PE-2 (nexthop).

```
A:PE-1# configure router mpls static-lsp "PE-1-PE-6"
A:PE-1>config>router>mpls>static-lsp$ to 192.0.2.6
A:PE-1>config>router>mpls>static-lsp$ push 100 nexthop 192.168.1.2
A:PE-1>config>router>mpls>static-lsp$ no shutdown
A:PE-1>config>router>mpls>static-lsp$ exit all
```

The transit LSRs (PE-2 and PE-5) perform swap operations and forward the packet to the manually defined next-hop. On the LSR under the context of the interface on which the incoming LSP arrives, the correct label is selected (label-map) and in this context a swap operation with a new label and the new next hop (nexthop) is entered.

```
A:PE-2>config>router>mpls# interface int-PE-2-PE-1
A:PE-2>config>router>mpls>if# label-map 100
A:PE-2>config>router>mpls>if>label-map$ swap 150 nexthop 192.168.4.2
A:PE-2>config>router>mpls>if>label-map$ no shutdown
A:PE-2>config>router>mpls>if>label-map$ exit all

A:PE-5# configure router mpls interface int-PE-5-PE-2
A:PE-5>config>router>mpls>if$ label-map 150
```

## Manually Configured LSPs

```
A:PE-5>config>router>mpls>if>label-map$ swap 200 nexthop 192.168.6.2
A:PE-5>config>router>mpls>if>label-map$ no shutdown
A:PE-5>config>router>mpls>if>label-map$ back
```

The terminating router PE-6 performs a POP operation and forwards the now unlabeled packets external to the MPLS domain.

```
A:PE-6>config>router>mpls# interface int-PE-6-PE-5
A:PE-6>config>router>mpls>if$ label-map 200
A:PE-6>config>router>mpls>if>label-map$ pop
A:PE-6>config>router>mpls>if>label-map$ no shutdown
A:PE-6>config>router>mpls>if>label-map$ exit all
```

To verify the operational status of the static LSP configuration the **show router mpls static-lsp** command is used on the iLER. A static LSP is considered to be operationally up when only its next-hop is reachable. Since there is no check whether the end-to-end LSP path is up (the LSP connectivity to the eLER is never verified), it can be the static LSP path is broken while the iLER displays an operational enabled LSP.

```
A:PE-1# show router mpls static-lsp
=====
MPLS Static LSPs (Originating)
=====
LSP Name      To                Next Hop          Out Label Up/Down Time  Adm  Opr
  ID                                     Out Port
-----
PE-1-PE-6    192.0.2.6         192.168.1.2      100      0d 00:01:50  Up   Up
  2                                     1/1/1
-----
LSPs : 1
=====
A:PE-1#
```

On the LSR the **transit** keyword is added to the command

```
A:PE-2# show router mpls static-lsp transit
=====
MPLS Static LSPs (Transit)
=====
In Label      In Port          Out Label  Out Port      Next Hop          Adm  Opr
-----
100           1/1/1           150        1/1/3         192.168.4.2      Up   Up
-----
LSPs : 1
=====
A:PE-2#
```

```
A:PE-5# show router mpls static-lsp transit
=====
MPLS Static LSPs (Transit)
```

```

=====
In Label    In Port    Out Label  Out Port   Next Hop           Adm  Opr
-----
150         1/1/1      200        1/1/3      192.168.6.2       Up   Up
-----
LSPs : 1
=====
A:PE-5#

```

And, on the terminating router (eLER), the keyword **terminate** is added.

```

A:PE-6# show router mpls static-lsp terminate
=====
MPLS Static LSPs (Terminate)
=====
In Label    In Port    Out Label  Out Port   Next Hop           Adm  Opr
-----
200         1/1/2      n/a        n/a        n/a                Up   Up
-----
LSPs : 1
=====
A:PE-6#

```

## Manually Configured LSPs

To track the label action associated with the static LSP configuration, the **show router mpls interface label-map** command can be used on all LSRs and eLERs (not iLER).

```
A:PE-2# show router mpls interface label-map
=====
MPLS Interfaces (Label-Map)
=====
In Label  In I/F      Out Label  Out I/F    Next Hop      Type      Adm  Opr
-----
100       1/1/1      150       1/1/3     192.168.4.2  Static    Up   Up
-----
Interfaces : 1
=====
A:PE-2#
```

```
A:PE-6# show router mpls interface label-map
=====
MPLS Interfaces (Label-Map)
=====
In Label  In I/F      Out Label  Out I/F    Next Hop      Type      Adm  Opr
-----
200       1/1/2      n/a       n/a       n/a           Static    Up   Up
-----
Interfaces : 1
=====
A:PE-6#
```

The **show router mpls status** command is used to verify each of the LSP types, the number configured and whether they originate on, transit through or terminate on the router.

```
A:PE-1# show router mpls status
=====
MPLS Status
=====
Admin Status      : Up           Oper Status      : Up
Oper Down Reason  : n/a
FR Object         : Enabled      Resignal Timer   : Disabled
Hold Timer        : 1 seconds   Next Resignal    : N/A
Srlg Frr          : Disabled     Srlg Frr Strict  : Disabled
Dynamic Bypass    : Enabled      User Srlg Database : Disabled
Least Fill Min Thd.: 5 percent  LeastFill ReoptiThd: 10 percent

Sec FastRetryTimer : Disabled

LSP Counts      Originate      Transit      Terminate
-----
Static LSPs     1              0            0
Dynamic LSPs    0              0            0
Detour LSPs     0              0            0
=====
A:PE-1#
```

PHP can be used with static LSPs. This is achieved by configuring the penultimate LER to swap the incoming label to implicit-null instead of a specific label value (the label-map must be shutdown to add the **swap** command).

```
A:PE-5# configure router mpls interface int-PE-5-PE-2
A:PE-5>config>router>mpls>if$ label-map 150
A:PE-5>config>router>mpls>if>label-map$ swap implicit-null-label nexthop 192.168.6.2
A:PE-5>config>router>mpls>if>label-map$ no shutdown
A:PE-5>config>router>mpls>if>label-map$ back
```

The previous configuration will cause PE-5 to pop the top label from the incoming labeled frame received from PE-2 and send it to PE-6 without adding another outer label. The result can be seen from the following command (note that label 3 is never actually pushed onto a frame).

```
A:PE-5# show router mpls static-lsp transit
=====
MPLS Static LSPs (Transit)
=====
In Label      In Port      Out Label    Out Port     Next Hop      Adm   Opr
-----
150           1/1/1        3            1/1/3        192.168.6.2  Up    Up
-----
LSPs : 1
=====
```

If the traffic arriving at PE-5 was IP with a single label then it would arrive at PE-6 as unlabeled IP traffic.

If the static LSP spans a single hop (PE-1 to PE-2) the ingress LER would push the implicit-null instead of using a swap.

```
A:PE-1# configure router mpls static-lsp PE-1-PE-2
A:PE-1>config>router>mpls>static-lsp$ to 192.0.2.2
A:PE-1>config>router>mpls>static-lsp$ push implicit-null-label nexthop 192.168.1.2
A:PE-1>config>router>mpls>static-lsp$ no shutdown
A:PE-1>config>router>mpls>static-lsp$ exit all
```

In this case, no MPLS action (swap or pop) is required for this LSP on PE-2 with respect to this LSP.

## LDP

LDP is a simple label distribution protocol with basic MPLS functionality (no traffic engineering and usually no recovery options). It relies on the underlying routing information provided by an IGP in order to forward labeled packets. Each LDP configured LSR will originate a label for its system address and a label for each FEC for which it has a next-hop that is external to the MPLS domain, without the explicit need to create the LSPs. When deviations from this default behavior are desired, import and export policies can be applied.

The configuration is as simple as enabling the LDP protocol instance and adding all network interfaces, for each node. As an example the configuration on node PE-1 is displayed below; similar configurations apply on the other nodes.

```
A:PE-1# configure router ldp
A:PE-1>config>router>ldp$ interface-parameters
A:PE-1>config>router>ldp>if-params$ interface int-PE-1-PE-2
A:PE-1>config>router>ldp>if-params>if$ back
A:PE-1>config>router>ldp>if-params# interface int-PE-1-PE-4
A:PE-1>config>router>ldp>if-params>if$ back
A:PE-1>config>router>ldp>if-params# back
```

The **show router ldp discover** and **show router ldp session** commands can be used to verify the LDP peers and sessions. The adjacency type (AdjType) needs to be **Link** while the state should be **Estab(lished)**.

```
A:PE-1# show router ldp discovery
=====
LDP Hello Adjacencies
=====
Interface Name                Local Addr    Peer Addr    AdjType State
-----
int-PE-1-PE-2                 192.0.2.1    192.0.2.2    Link   Estab
int-PE-1-PE-4                 192.0.2.1    192.0.2.4    Link   Estab
=====
No. of Hello Adjacencies: 2
=====
A:PE-1#
```

```
A:PE-1# show router ldp session
=====
LDP Sessions
=====
Peer LDP Id      Adj Type  State        Msg Sent  Msg Recv  Up Time
-----
192.0.2.2:0     Link     Established  157       158       0d 00:06:56
192.0.2.4:0     Link     Established  33        34        0d 00:01:12
=====
No. of Sessions: 2
=====
```



## LDP

192.0.2.3/32	Swap	131068	131067	1/1/1	192.168.1.2
192.0.2.4/32	Push	--	131065	1/1/2	192.168.2.2
192.0.2.4/32	Swap	131065	131065	1/1/2	192.168.2.2
192.0.2.5/32	Push	--	131069	1/1/1	192.168.1.2
192.0.2.5/32	Swap	131069	131069	1/1/1	192.168.1.2
192.0.2.6/32	Push	--	131068	1/1/1	192.168.1.2
192.0.2.6/32	Swap	131067	131068	1/1/1	192.168.1.2

-----  
No. of Prefix Bindings: 11  
=====

A:PE-1#

In order to signal PHP with LDP, implicit-null must be configured on the eLER.

A:PE-6# configure router ldp implicit-null-label

The implicit-null is signaled immediately, all related labels are withdrawn and re-advertised with the label 3. The new label would show up on PE-5 as a swap from the ingress label to an egress label of 3. If the traffic arriving at PE-5 was IP with a single label then it would arrive at PE-6 as unlabeled IP traffic.



## Import and Export Policies

The default label handling behavior is to originate label bindings for the system address and to propagate all FECs received. If this is not the desired behavior, an import/export policy can be applied. An export policy may be configured to control the set of LDP label bindings advertised by the LER (sending to LDP peers). As such, export policies are used to include additional FECs rather than filtering FECs from those advertised. An import policy can be used to control for which FECs a router will generate labels (accepting from LDP peers). This functionality is not unique to LDP; it can be used for RSVP-TE, OSPF, and IS-IS as well as others.

By default no labels are generated for directly connected (local) interfaces. To change this behavior, an export policy is created and applied to the LDP instance. There is no configuration difference in defining an import and export policy.

A policy starts with the keyword **begin** contains a list of entries (of which each has a number), and ends with the keyword **commit**. An entry typically contains matching criteria (however, it is not required in such cases where everything matches) and a corresponding action. Entries without an action are considered incomplete and are rendered inactive. When executing the policy, the router executes the specified action on the first matching statement; it does not process any further matches. For this reason, entries must be sequenced correctly from most to least specific.

The configuration of the LDP export policy for local interfaces is given below.

```
A:PE-1# configure router policy-options
A:PE-1>config>router>policy-options# begin
A:PE-1>config>router>policy-options# policy-statement LDP-export
A:PE-1>config>router>policy-options>policy-statement$ entry 10
A:PE-1>config>router>policy-options>policy-statement>entry$ from protocol direct
A:PE-1>config>router>policy-options>policy-statement>entry# action accept
A:PE-1>config>router>policy-options>policy-statement>entry>action# back
A:PE-1>config>router>policy-options>policy-statement>entry# back
A:PE-1>config>router>policy-options>policy-statement# back
A:PE-1>config>router>policy-options# commit
A:PE-1>config>router>policy-options# exit all
```

There are 11 active LDP bindings before applying the export policy.

```
A:PE-1# show router ldp bindings active
=====
Legend: (S) - Static
=====
LDP Prefix Bindings (Active)
=====
Prefix                Op   IngLbl   EgrLbl   EgrIntf/LspId  EgrNextHop
-----
192.0.2.1/32          Pop  131071   --        --              --
192.0.2.2/32          Push --        131071   1/1/1          192.168.1.2
192.0.2.2/32          Swap 131070   131071   1/1/1          192.168.1.2
192.0.2.3/32          Push --        131069   1/1/1          192.168.1.2
192.0.2.3/32          Swap 131069   131069   1/1/1          192.168.1.2
192.0.2.4/32          Push --        131065   1/1/2          192.168.2.2
192.0.2.4/32          Swap 131068   131065   1/1/2          192.168.2.2
192.0.2.5/32          Push --        131067   1/1/1          192.168.1.2
192.0.2.5/32          Swap 131067   131067   1/1/1          192.168.1.2
192.0.2.6/32          Push --        131066   1/1/1          192.168.1.2
192.0.2.6/32          Swap 131066   131066   1/1/1          192.168.1.2
-----
No. of Prefix Bindings: 11
=====
A:PE-1#
```

The LDP export or import policy is applied to the LDP instance on the router, respectively, with the **export** or **import** keyword.

```
A:PE-1# configure router ldp export LDP-export
```

When the export policy is applied, the active LDP binding table has two additional entries: the local interfaces of PE-1.

```
A:PE-1# show router ldp bindings active
=====
Legend: (S) - Static
=====
LDP Prefix Bindings (Active)
=====
Prefix                Op   IngLbl   EgrLbl   EgrIntf/LspId  EgrNextHop
-----
192.0.2.1/32          Pop  131071   --        --              --
192.0.2.2/32          Push --        131071   1/1/1          192.168.1.2
192.0.2.2/32          Swap 131070   131071   1/1/1          192.168.1.2
192.0.2.3/32          Push --        131069   1/1/1          192.168.1.2
192.0.2.3/32          Swap 131069   131069   1/1/1          192.168.1.2
192.0.2.4/32          Push --        131065   1/1/2          192.168.2.2
192.0.2.4/32          Swap 131068   131065   1/1/2          192.168.2.2
192.0.2.5/32          Push --        131067   1/1/1          192.168.1.2
192.0.2.5/32          Swap 131067   131067   1/1/1          192.168.1.2
```

## Point-to-Point LSPs

192.0.2.6/32	Push	--	131066	1/1/1	192.168.1.2
192.0.2.6/32	Swap	131066	131066	1/1/1	192.168.1.2
192.168.1.0/30	Pop	131065	--	--	--
192.168.2.0/30	Pop	131064	--	--	--

-----  
No. of Prefix Bindings: 13

=====  
A:PE-1#

## RSVP-TE

RSVP-TE, an extension of the original RSVP protocol, has two major benefits adding to the basic MPLS functionality. The first benefit is traffic engineering, which allows the ingress router to create an LSP with specific constraints beyond the best route chosen by the IGP. The second benefit is improved network resiliency and shorter service interruption times when a link or node fails in the network.

In this section, an RSVP-TE based LSP is established from PE-1 to PE-6, starting with a simple LSP with no specific TE constraints. Although the Fast Reroute (FRR), admin groups, hop limit restriction, and bandwidth reservation features could be implemented, the usage of these features is beyond the scope of this document.

Like the configuration of static LSPs, the MPLS instance needs to be enabled on each router and all network interfaces facing the MPLS domain as well as the system interface. When adding interfaces to the MPLS instance they are automatically added to the RSVP instance as well, but the instance itself is still in an administrative shutdown state. The next step is to enable the RSVP instance on all routers in the MPLS network. As a result all interfaces facing the MPLS domain as well as the system interface are added to the MPLS and RSVP instance and both instances are in a no shutdown state. For PE-1 this comes to:

```
A:PE-1# configure router rsvp
A:PE-1>config>router>rsvp# no shutdown
A:PE-1>config>router>rsvp# info
-----
          interface "system"
          exit
          interface "int-PE-1-PE-2"
          exit
          interface "int-PE-1-PE-4"
          exit
          no shutdown
-----
A:PE-1>config>router>rsvp#
```

**Strict or loose path** On the iLER first the definition of a path is required. A path is a sequence of MPLS routers (hops) through which the LSP -using that path- has to pass. It is not uniquely bound to a particular LSP; it can be used by any LSP originating in that node. A hop in a path can be strict or loose: strict or loose meaning that the LSP must take either a direct path from the previous hop router to this router (strict) or can traverse through other routers (loose). The hops missing in the loose path definition are created by calculating the IGP shortest path. A third possibility is an entirely empty path implying not a single node is required to be present in the LSP path and the shortest path from the IGP is used to define the LSP path. The use of other techniques, such as the use of admin groups or shared risk link groups, are not covered in this document. Three paths are configured below, respectively:

1. Only strict hops
7. Mixed strict and loose hops
8. Completely loose path

To find a valid path, the last hop in the path sequence needs to be the system IP or an interface address of the terminating router (eLER). The IP addresses in the hop command can be the node's system IP addresses or its interface addresses. However, it is recommended to use the system IP addresses as this allows more flexibility when finding new paths in failover scenarios (because the upstream node could use any of multiple paths to the system address, where specifying the interface address would restrict the upstream node to a single entry-point).

```
A:PE-1# configure router mpls path "strict_path"
A:PE-1>config>router>mpls>path$ hop 1 192.0.2.2 strict
A:PE-1>config>router>mpls>path$ hop 2 192.0.2.5 strict
A:PE-1>config>router>mpls>path$ hop 3 192.0.2.6 strict
A:PE-1>config>router>mpls>path$ no shutdown

A:PE-1# configure router mpls
A:PE-1>config>router>mpls# path "loose_path"
A:PE-1>config>router>mpls>path$ hop 1 192.0.2.5 loose
A:PE-1>config>router>mpls>path$ hop 2 192.0.2.6 strict
A:PE-1>config>router>mpls>path$ no shutdown

A:PE-1# configure router mpls path "completely_loose_path"
A:PE-1>config>router>mpls>path$ no shutdown
```

The paths can be checked with the **show router mpls path** command.

```
A:PE-1# show router mpls path
=====
MPLS Path:
=====
Path Name                               Adm  Hop Index  IP Address  Strict/Loose
-----
strict_path                               Up   1          192.0.2.2   Strict
                                           2          192.0.2.5   Strict
                                           3          192.0.2.6   Strict

loose_path                                Up   1          192.0.2.5   Loose
                                           2          192.0.2.6   Strict

completely_loose_path                     Up   no hops    n/a         n/a

-----
Paths : 3
=====
A:PE-1#
```

## Simple RSVP LSP

The configuration of a simple LSP using RSVP signaling contains at least on the iLER:

- System IP address of the terminating node (to)
- Path the LSP will take to the eLER (primary)
- Administratively enabled (no shutdown)

```
A:PE-1>configure router mpls lsp "lsp-PE-1-PE-6_simple"
A:PE-1>config>router>mpls>lsp# to 192.0.2.6
A:PE-1>config>router>mpls>lsp# primary "strict_path"
A:PE-1>config>router>mpls>lsp# no shutdown
```

The nodes through which the LSP will pass (LSRs and eLER) require no additional configuration: enabling MPLS (and automatically RSVP together with it) on their interfaces suffices.

An overview of all LSPs configured on a particular node is given by the **show router mpls lsp** command. More details about a particular LSP can be retrieved by adding the keyword **detail** to the previous command.

```
A:PE-1# show router mpls lsp
=====
MPLS LSPs (Originating)
=====
LSP Name                               To                               Fastfail   Adm   Opr
                               Config
-----
lsp-PE-1-PE-6_simple                   192.0.2.6                       No         Up    Up
-----
LSPs : 1
=====
A:PE-1#
```

```
A:PE-1# show router mpls lsp lsp-PE-1-PE-6_simple detail
=====
MPLS LSPs (Originating) (Detail)
=====
Type : Originating
-----
LSP Name      : lsp-PE-1-PE-6_simple           LSP Tunnel ID : 4
From          : 192.0.2.1                       To            : 192.0.2.6
Adm State     : Up                           Oper State    : Up
LSP Up Time  : 0d 00:07:33                    LSP Down Time : 0d 00:00:00
Transitions  : 5                             Path Changes  : 5
Retry Limit   : 0                             Retry Timer   : 30 sec
Signaling     : RSVP                         Resv. Style   : SE
Hop Limit    : 255                           Negotiated MTU : 1564
Adaptive     : Enabled                       ClassType     : 0
FastReroute  : Disabled                      Oper FR       : Disabled
CSPF         : Disabled                      ADSPEC        : Disabled
```

```

Metric      : 0
Include Grps:          Exclude Grps :
None        :          None         :
Type        : RegularLsp           Least Fill  : Disabled
LdpOverRsvp : Enabled              VprnAutoBind : Enabled
Oper Metric : 65535

Primary(a)  : strict_path           Up Time     : 0d 00:07:33
Bandwidth   : 0 Mbps

=====
A:PE-1#

```

In order to signal PHP with RSVP, implicit-null must be configured on the eLER (RSVP must be shutdown to perform this command).

```
A:PE-6# configure router rsvp implicit-null-label
```

The implicit-null is signaled after re-enabling RSVP and would show up on PE-5 as an egress label of 3.

The use of implicit-null can also be enabled/disabled on a per interface basis (either RSVP, or the interface within RSVP, must be shutdown to perform this change).

```

A:PE-6>config>router>rsvp# interface int-PE-6-PE-5
A:PE-6>config>router>rsvp>if# implicit-null-label
  - implicit-null-label {<enable|disable>}
  - no implicit-null-label

```

If the traffic arriving at PE-5 was IP with a single label then it would arrive at PE-6 as unlabeled IP traffic.

---

## Manual Resignal

Instead of waiting for the resignal timer to expire, one can manually trigger the resignal process. The command includes both the LSP name and path to resignal:

```
A:PE-1# tools perform router mpls resignal lsp lsp-PE-1-PE-6_simple path strict_path
```

## LSP OAM

The LSP diagnostics are modeled after ICMP echo request/reply which provides a mechanism to detect data plane failures in MPLS LSPs. For a given FEC, LSP ping verifies whether the packet reaches the egress label edge router (LER). While in LSP traceroute mode, the packet is sent to the control plane of each transit label switched router (LSR) which performs various checks to see if it is actually a transit LSR for the path.

```
A:PE-1# oam lsp-ping lsp-PE-1-PE-6_simple
LSP-PING lsp-PE-1-PE-6_simple: 92 bytes MPLS payload
Seq=1, send from intf int-PE-1-PE-2, reply from 192.0.2.6
      udp-data-len=32 ttl=255 rtt=3.21ms rc=3 (EgressRtr)

---- LSP lsp-PE-1-PE-6_simple PING Statistics ----
1 packets sent, 1 packets received, 0.00% packet loss
round-trip min = 3.21ms, avg = 3.21ms, max = 3.21ms, stddev = 0.000ms

A:PE-1# oam lsp-trace lsp-PE-1-PE-6_simple
lsp-trace to lsp-PE-1-PE-6_simple: 0 hops min, 0 hops max, 116 byte packets
1 192.0.2.2 rtt=11.9ms rc=8 (DSRtrMatchLabel)
2 192.0.2.3 rtt=6.35ms rc=8 (DSRtrMatchLabel)
3 192.0.2.6 rtt=4.85ms rc=3 (EgressRtr)
```



## Debug Tools

A wide range of debug tools are available which can be tuned to the specific information of importance for a certain troubleshooting task. In the **debug router mpls** context, the LSP object to trace or monitor can be selected by the following parameters:

- LSP name
- Source address of the LSP (the **from** parameter in the LSP definition)
- Termination point of the LSP (the **to** parameter in the LSP definition)
- Tunnel ID of the LSP
- LSP ID

```
A:PE-1# debug router mpls
- mpls [lsp <lsp-name>] [sender <source-address>] [endpoint
  <endpoint-address>] [tunnel-id <tunnel-id>] [lsp-id <lsp-id>]
- no mpls

<lsp-name>          : [80 chars max]
<source-address>   : a.b.c.d
<endpoint-address> : a.b.c.d
<tunnel-id>        : [0..4294967295]
<lsp-id>           : [1..65535]

[no] event          + Enable/disable debugging for specific MPLS events
```

In the **debug** command tree, the MPLS event type can be selected (tracing must be enabled):

```
A:PE-1# debug router mpls lsp lsp-PE-1-PE-6_simple event
- event
- no event

[no] all           - Enable/disable debugging for MPLS all
[no] frr           - Enable/disable debugging for MPLS frr
[no] iom           - Enable/disable debugging for MPLS iom
[no] lsp-setup     - Enable/disable debugging for MPLS lsp setup
[no] mbb           - Enable/disable debugging for MPLS mbb
[no] misc          - Enable/disable debugging for MPLS misc
[no] xc            - Enable/disable debugging for MPLS xc
```

As an example, the **all** keyword is entered, logging all MPLS events related to the selected LSP:

```
A:PE-1# debug router mpls lsp lsp-PE-1-PE-6_simple event all
```

The last step is to create a log container which will gather all MPLS debugging information according to the criteria set in the debug context. The **from debug-trace** parameter must be configured but there are several options where the different captured entries will be stored: console, a syslog server, SNMP, local file on the compact flash card, a temporary circular memory buffer, or the telnet/SSH session from which you are logged into the node.

The log container's ID is just a local number without any other significance.

```
A:PE-1# configure log log-id 2
A:PE-1>config>log>log-id# from debug-trace
A:PE-1>config>log>log-id# to
  - to console
  - to file <log-file-id>
  - to memory [<size>]
  - to session
  - to snmp [<size>]
  - to syslog <syslog-id>

<console>          : keyword - specifies console as destination
<syslog-id>        : [1..10]
<snmp>             : keyword - specifies SNMP as destination
<log-file-id>      : [1..99]
<memory>           : keyword - specifies memory as destination
<size>             : [50..1024]
<session>          : keyword - specifies telnet session as destination
```

For this example, the temporary buffer (with adjustable size) is chosen and the log container is enabled (no shutdown).

```
A:PE-1>config>log>log-id# to memory
A:PE-1>config>log>log-id# no shutdown
```

All MPLS events related to the selected LSP are stored in the location (memory) specified. The content of this log container can be viewed through the **show log log-id 2** command.

```
A:PE-1# show log log-id 2
=====
Event Log 2
=====
Description : (Not Specified)
Memory Log contents [size=100  next event=13  (not wrapped)]

12 2010/01/29 17:30:06.78 UTC MINOR: DEBUG #2001 Base MPLS
"MPLS: RTM
Update traffic-engineering information for interface index 2
Oper state Up, Admin group bitmap 0x0, Max bandwidth 125000000Bps
Max reservable bandwidth 125000000Bps, Available bandwidth 0Bps"

11 2010/01/29 17:30:06.77 UTC MINOR: DEBUG #2001 Base MPLS
```

```
"MPLS: ICC  
Set protect group 127, protect instance 3, use protected yes"  
  
10 2010/01/29 17:30:00.82 UTC MINOR: DEBUG #2001 Base MPLS  
"MPLS: RTM  
Process RT Internal Unreg event for Pref 192.0.2.2/32"  
  
...
```

## Conclusion

MPLS provides the capability to establish connection oriented paths over a connectionless network. The LSP offers a mechanism to engineer network traffic on constraint-based paths rather than the IGP shortest path. This can greatly improve network resiliency. In this section, the configuration of several LSP features is given together with the associated show output which can be used to verify and troubleshoot.