

# MPLS Traffic Engineering

MPLS TE, FRR, and Merge Point concepts + Configuration



<b>Introduction</b> .....	<b>3</b>
<b>1. MPLS TE Configuration</b> .....	<b>4</b>
1.1 PE1 Configuration .....	<b>4</b>
1.2 P1, P2, P3 and P4 Configuration .....	<b>5</b>
1.3 PE2 Configuration .....	<b>6</b>
1.4 Verification .....	<b>6</b>
<b>2. MPLS TE Fast Reroute Link Protection</b> .....	<b>9</b>
2.1 Backup Tunnel Configuration on P1 .....	<b>9</b>
2.2 Backup Tunnel Configuration on P3 .....	<b>10</b>
2.3 PE1 Configuration .....	<b>10</b>
2.4 Verification .....	<b>10</b>
<b>3. Testing MPLS TE Fast Reroute Link Protection</b> .....	<b>11</b>
<b>Conclusion</b> .....	<b>13</b>

## Introduction

In MPLS Traffic Engineering (TE), a Merge Point (MP) is a router or a point in the network where backup paths merge with primary paths to provide faster traffic restoration in case of link or node failures.

Fast Reroute (FRR) is a technique used in MPLS TE to minimize the impact of link or node failures on traffic. When a failure occurs, FRR allows traffic to be quickly rerouted onto a backup path without waiting for the network to converge on a new path. Therefore, MP is a critical component of FRR because it is where the backup and primary paths converge.

When a failure occurs on the primary path, the backup path is activated, and traffic is rerouted to the MP. At the Merge Point, the traffic from the primary and backup paths are merged and forwarded toward the destination.

To illustrate this concept, consider the topology depicted in Figure 1. In the topology, the red path represents the primary LSP (Label Switched Path) and carries the traffic. The green path, on the other hand, represents the backup LSP and is used only in case of failure. If a link fails, such as between Router P1 and P3, the traffic cannot be carried between Router PE1 and PE2 through the red LSP.




Figure 1 - Network Topology

A fast reroute is used to ensure that traffic can continue to flow over an alternate path within 50 milliseconds. When the link between P1 and P3 fails, Router P1 learns about the failure first, since it is directly attached to the failed link. Router P1 then redirects traffic to the green LSP, which is terminated at router P3. The node P3 is called the Merge Point (MP) because the backup LSP is merged with the primary LSP at that node. If traffic reaches Node P1 after the P1-P3 link failure, Node P3 can continue forwarding traffic to the destination (Router PE2) over the green LSP.

## 1. MPLS TE Configuration

MPLS tunnels are unidirectional, which means we need one tunnel in each direction to send traffic. In the case of tunnel12 from PE1 to PE2, the headend or start of the tunnel is Router PE1 with 20.20.20.1 configured, while the tailend is 50.50.50.2 configured on PE2. Similarly, for tunnel 21 from PE2 to PE1, the headend is Router PE2, and the tailend is PE1.

To exchange link and loopback subnets, OSPF is enabled on all routers.

To reserve and exchange labels, it is necessary to enable RSVP on all interfaces except loopback interfaces.

### 1.1 PE1 Configuration

```
mpls traffic-eng tunnels
```

```
interface Loopback0
ip address 2.2.2.2 255.255.255.255
```

Note that the Loopback1 IP address 192.168.1.1 is used for testing tunnel connectivity between PE1 and PE2 routers and is not related to the MPLS TE configuration.

```
interface Loopback1
ip address 192.168.1.1 255.255.255.0
```

Enable MPLS TE and RSVP on all Gigabit-Ethernet interfaces with a bandwidth reservation of 512 kbps.

```
interface GigabitEthernet0/0
ip address 20.20.20.1 255.255.255.252
mpls traffic-eng tunnels
ip rsvp bandwidth 1024 1024
```

Enable MPLS TE and RSVP on all Gigabit-Ethernet interfaces with a bandwidth reservation of 512 kbps.

```
router ospf 1
mpls traffic-eng router-id Loopback0
mpls traffic-eng area 0
router-id 2.2.2.2
network 2.2.2.2 0.0.0.0 area 0
network 20.20.20.0 0.0.0.3 area 0
```

# MPLS Traffic Engineering

---

Define explicit path for PE1-PE2 tunnel: PE1-> P1->P3-> P4 and PE2.

```
ip explicit-path name PE1-PE2 enable
  next-address 20.20.20.2
  next-address 30.30.30.6
  next-address 40.40.40.6
  next-address 50.50.50.2
```

Create a tunnel interface, and define the MPLS-TE tunnel encapsulation method, tunnel destination, bandwidth, priority, and other relevant parameters.

```
interface Tunnel12
  ip unnumbered Loopback0
  tunnel mode mpls traffic-eng
  tunnel destination 8.8.8.8
  tunnel mpls traffic-eng priority 2 2
  tunnel mpls traffic-eng bandwidth 512
  tunnel mpls traffic-eng path-option 1 explicit name PE1-PE2
  tunnel mpls traffic-eng fast-reroute
  no routing dynamic
```

We also need to route traffic to the 192.168.2.1 loopback interface of PE2 through the tunnel12 interface.

```
ip route 192.168.2.0 255.255.255.0 Tunnel12
```

## 1.2 P1, P2, P3 and P4 Configuration

We will only showcase the configuration of the P1 transit router, as the configuration of the other routers is similar and, therefore, not necessary to present.

```
mpls traffic-eng tunnels

interface Loopback0
  ip address 4.4.4.4 255.255.255.255

interface GigabitEthernet0/0
  ip address 20.20.20.2 255.255.255.252
  mpls traffic-eng tunnels
  ip rsvp bandwidth 1024 1024

interface GigabitEthernet0/1
  ip address 30.30.30.5 255.255.255.252
  mpls traffic-eng tunnels
  ip rsvp bandwidth 1024 1024

interface GigabitEthernet0/2
  ip address 30.30.30.1 255.255.255.252
  mpls traffic-eng tunnels
  ip rsvp bandwidth 1024 1024

router ospf 1
  mpls traffic-eng router-id Loopback0
  mpls traffic-eng area 0
  router-id 4.4.4.4
  network 0.0.0.0 255.255.255.255 area 0
```

## 1.3 PE2 Configuration

```
mpls traffic-eng tunnels

interface Loopback0
 ip address 8.8.8.8 255.255.255.255

interface Loopback1
 ip address 192.168.2.1 255.255.255.0

interface GigabitEthernet0/0
 ip address 50.50.50.2 255.255.255.252
 mpls traffic-eng tunnels
 ip rsvp bandwidth 1024 1024

router ospf 1
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng area 0
 router-id 8.8.8.8
 network 8.8.8.8 0.0.0.0 area 0
 network 50.50.50.0 0.0.0.3 area 0

ip explicit-path name PE2-PE1 enable
 next-address 50.50.50.1
 next-address 40.40.40.1
 next-address 30.30.30.1
 next-address 20.20.20.1

interface Tunnel121
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 2.2.2.2
```

```
tunnel mpls traffic-eng priority 2 2
 tunnel mpls traffic-eng bandwidth 512
 tunnel mpls traffic-eng path-option 1 explicit name PE2-PE1
 tunnel mpls traffic-eng fast-reroute
 no routing dynamic
```

```
ip route 192.168.1.0 255.255.255.0 Tunnel121
```

## 1.4 Verification

To begin, we verify the operational status of tunnel 12 on router PE1, as shown in Figure 2.

```
PE1# show mpls traffic-eng tunnels
```

```
Name: PE1_t12 (Tunnel12) Destination: 8.8.8.8
Status:
Admin: up Oper: up Path: valid Signalling: connected
path option 1, type explicit PE1-PE2 (Basis for Setup, path weight 4)

Config Parameters:
Bandwidth: 512 kbps (Global) Priority: 2 2 Affinity: 0x0/0xFFFF
Metric Type: TE (default)
AutoRoute: disabled LockDown: disabled Loadshare: 512 bw-based
auto-bw: disabled
Active Path Option Parameters:
State: explicit path option 1 is active
BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
```

# MPLS Traffic Engineering

```
InLabel : -
OutLabel : GigabitEthernet0/0, 17
RSVP Signalling Info:
  Src 2.2.2.2, Dst 8.8.8.8, Tun_Id 12, Tun_Instance 15
RSVP Path Info:
  My Address: 20.20.20.1
  Explicit Route: 20.20.20.2 30.30.30.5 30.30.30.6 40.40.40.5
                  40.40.40.6 50.50.50.1 50.50.50.2 8.8.8.8
  Record Route: NONE
  Tspec: ave rate=512 kbits, burst=1000 bytes, peak rate=512 kbits
RSVP Resv Info:
  Record Route: 4.4.4.4(17) 6.6.6.6(16)
                7.7.7.7(17) 8.8.8.8(0)
  Fspec: ave rate=512 kbits, burst=1000 bytes, peak rate=512 kbits
Shortest Unconstrained Path Info:
  Path Weight: 4 (TE)
  Explicit Route: 20.20.20.1 20.20.20.2 30.30.30.1 30.30.30.2
                  40.40.40.1 40.40.40.2 50.50.50.1 50.50.50.2
                  8.8.8.8
History:
Tunnel:
  Time since created: 1 hours, 47 minutes
  Time since path change: 1 hours, 47 minutes
  Number of LSP IDs (Tun_Instances) used: 15
Current LSP:
  Uptime: 1 hours, 47 minutes
```

Figure 2 - Checking MPLS TE Tunnels on PE1 - Tunnel12

Figure 2 shows the RSVP Resv path label information used by each hop along the path from PE1 to PE2. At the start of the path, PE1 (2.2.2.2) pushes label 17 to the packets and forwards them to P1. Upon receiving the packets, P1 (4.4.4.4) swaps label 17 with a new label 16 before forwarding the traffic to P3, as shown in Figure 3. Then, at P3 (6.6.6.6), the incoming label 16 is swapped with a new label 17, and the traffic is forwarded to P4. Finally, at PE4 (7.7.7.7), label 17 is popped, and the traffic is sent to PE2 with no MPLS header.

```
P1#show mpls forwarding-table
Local   Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label   Label     or Tunnel Id   Switched     interface
16      Pop Label  8.8.8.8 21 [11]        0            Gi0/0       20.20.20.1
17      16        2.2.2.2 12 [15]       1180         Gi0/1       30.30.30.6
P1#
```

Figure 3 - P1 MPLS Forwarding Table

We can see in Figure 4 that the MPLS TE Tunnel31 on PE2 is also up and running. Any traffic destined for PE2 is placed into the tunnel and takes a path that passes through P4, P2, P1, and finally reaches PE1.

```
Name: PE2_t21 (Tunnel21) Destination: 2.2.2.2
Status:
Admin: up Oper: up Path: valid Signalling: connected
path option 1, type explicit PE2-PE1 (Basis for Setup, path weight 4)

Config Parameters:
Bandwidth: 512 kbps (Global) Priority: 2 2 Affinity: 0x0/0xFFFF
Metric Type: TE (default)
AutoRoute: disabled LockDown: disabled Loadshare: 512 bw-based
auto-bw: disabled
Active Path Option Parameters:
State: explicit path option 1 is active
BandwidthOverride: disabled LockDown: disabled Verbatim: disabled
```



# MPLS Traffic Engineering

```
InLabel : -
OutLabel : GigabitEthernet0/0, 16
RSVP Signalling Info:
  Src 8.8.8.8, Dst 2.2.2.2, Tun_Id 21, Tun_Instance 11
RSVP Path Info:
  My Address: 50.50.50.2
  Explicit Route: 50.50.50.1 40.40.40.2 40.40.40.1 30.30.30.2
                  30.30.30.1 20.20.20.2 20.20.20.1 2.2.2.2
  Record Route: NONE
  Tspec: ave rate=512 kbits, burst=1000 bytes, peak rate=512 kbits
RSVP Resv Info:
  Record Route: 7.7.7.7(16) 5.5.5.5(16)
                4.4.4.4(16) 2.2.2.2(0)
  Fspec: ave rate=512 kbits, burst=1000 bytes, peak rate=512 kbits
History:
Tunnel:
  Time since created: 3 hours, 43 minutes
  Time since path change: 3 hours, 42 minutes
  Number of LSP IDs (Tun_Instances) used: 11
Current LSP:
  Uptime: 3 hours, 42 minutes
```

```
LSP Tunnel PE1_t12 is signalled, connection is up
InLabel : GigabitEthernet0/0, implicit-null
OutLabel : -
RSVP Signalling Info:
  Src 2.2.2.2, Dst 8.8.8.8, Tun_Id 12, Tun_Instance 15
RSVP Path Info:
  My Address: 8.8.8.8
  Explicit Route: NONE
  Record Route: NONE
  Tspec: ave rate=512 kbits, burst=1000 bytes, peak rate=512 kbits
RSVP Resv Info:
  Record Route: NONE
  Fspec: ave rate=512 kbits, burst=1000 bytes, peak rate=512 kbits
```

Figure 4 - Checking MPLS TE Tunnels on PE2 - Tunnel 21

To check connectivity and make sure that traffic is passing through the tunnel, we use the traceroute command:

```
PE1#traceroute 192.168.2.1
Type escape sequence to abort.
Tracing the route to 192.168.2.1
VRF info: (vrf in name/id, vrf out name/id)
 1 20.20.20.2 [MPLS: Label 17 Exp 0] 12 msec 6 msec 4 msec
 2 30.30.30.6 [MPLS: Label 16 Exp 0] 8 msec 10 msec 13 msec
 3 40.40.40.6 [MPLS: Label 17 Exp 0] 6 msec 3 msec 15 msec
 4 50.50.50.2 15 msec 11 msec *
PE1#
```

Figure 5 - Checking Connectivity from PE1 to PE2



## 2. MPLS TE Fast Reroute Link Protection

We have set up a primary MPLS TE tunnel between PE1 and PE2, with a path through P1, P3, and P4. However, if the link between P1 and P3 fails, we'll lose connectivity to the tunnel endpoint, since the path is explicitly configured on PE1.

MPLS TE Fast Reroute (FRR) is a mechanism that provides fast protection against link and node failures in the network, ensuring that traffic is rerouted quickly and efficiently to avoid data loss. By default, MPLS TE calculates a new best path when there is a failure in the primary tunnel, but this process can be too slow for delay-sensitive applications such as VOIP. With MPLS TE fast reroute, traffic is forwarded down a backup tunnel when there is a failure in the primary tunnel, allowing a faster recovery time of less than 50 ms.

To bypass the link between P1 and P3, we will configure a backup NHOP (Next Hop) tunnel from P1 through P2, terminating at P3. The headend router of the backup tunnel is node P1, and the tailend router is the merge point (MP) - node P3. If the link between P1 and P3 fails, we'll still have a path to the TE tunnel endpoint, avoiding data loss.

In the following parts, we will enable link protection backup tunnel 100 on both P1 and P3 and associate the interface Gi0/1 with the backup tunnel on P1. When the interface Gi0/1 goes down, the backup tunnel is activated.

### 2.1 Backup Tunnel Configuration on P1

Assuming that the primary tunnel from PE1 to P2 is already configured and operational, we start configuring the backup tunnel on P1.

```
interface Tunnel100
  description "Link Protection Tunnel"
  ip unnumbered Loopback0
  tunnel mode mpls traffic-eng
  tunnel destination 6.6.6.6
  tunnel mpls traffic-eng path-option 1 explicit name P1-P3
  no routing dynamic

interface GigabitEthernet0/1
  ip address 30.30.30.5 255.255.255.252
  mpls traffic-eng tunnels
  mpls traffic-eng backup-path Tunnel100
  ip rsvp bandwidth 1024 1024

ip explicit-path name P1-P3 enable
next-address 30.30.30.2
next-address 40.40.40.10
```

## 2.2 Backup Tunnel Configuration on P3

```
interface Tunnel100
  description "Protection tunnel"
  ip unnumbered Loopback0
  tunnel mode mpls traffic-eng
  tunnel destination 4.4.4.4
  tunnel mpls traffic-eng path-option 1 explicit name P3-P1
  no routing dynamic

ip explicit-path name P3-P1 enable
  next-address 40.40.40.9
  next-address 30.30.30.1
```

## 2.3 PE1 Configuration

After configuring a backup tunnel on both P1 and P3 and associating it with the gi0/1 on both P1, one more step is required to ensure its proper functioning. Specifically, fast-reroute must be enabled on the TE tunnel configured on PE1 so that tunnel12 is fast-rerouted.

```
interface Tunnel12
  tunnel mpls traffic-eng fast-reroute
```

## 2.4 Verification

Firstly, we will check the status of the backup tunnel on P1 router (Figure 6):

```
P1#show mpls traffic-eng fast-reroute database
Headend frr information:
Protected tunnel          In-label Out intf/label  FRR intf/label  Status

LSP midpoint frr information:
LSP identifier           In-label Out intf/label  FRR intf/label  Status
2.2.2.2 12 [15]         17      Gi0/1:16       Tu100:16       ready
P1#
```

Figure 6 - Checking MPLS TE FRR Database Before Link Failure Between P1 and P3

The backup tunnel is currently in a ready state and will be activated automatically if interface Gi0/1 on P1 goes down. When packets are sent from PE1 to P1, they are labeled 17 and forwarded with label 16. The router P3 expects the label to be 16, but if the link between P1 and P3 path to label 16 is broken, we must ensure that label 16 remains protected for proper packet forwarding.

Therefore, packets sent from P1 to P2 via the backup path from the Gi0/2 interface when the backup tunnel is activated will have two labels. The inner label expected by P3 is 16, and the outer label inserted by P1 is 18 (Figure 7).

# MPLS Traffic Engineering

```
P1#show mpls traffic-eng tunnels tunnel 100
```

```
Name: "Link Protection Tunnel"          (Tunnel100) Destination: 6.6.6.6
Status:
  Admin: up          Oper: up          Path: valid          Signalling: connected
  path option 1, type explicit P1-P3 (Basis for Setup, path weight 2)

Config Parameters:
  Bandwidth: 0          kbps (Global) Priority: 7 7  Affinity: 0x0/0xFFFF
  Metric Type: TE (default)
  AutoRoute: disabled LockDown: disabled Loadshare: 0          bw-based
  auto-bw: disabled

Active Path Option Parameters:
  State: explicit path option 1 is active
  BandwidthOverride: disabled LockDown: disabled Verbatim: disabled

InLabel : -
OutLabel : GigabitEthernet0/2, 18
RSVP Signalling Info:
  Src 4.4.4.4, Dst 6.6.6.6, Tun_Id 100, Tun_Instance 14
RSVP Path Info:
  My Address: 30.30.30.1
  Explicit Route: 30.30.30.2 40.40.40.9 40.40.40.10 6.6.6.6
  Record Route: NONE
  Tspec: ave rate=0 kbits, burst=1000 bytes, peak rate=0 kbits
RSVP Resv Info:
  Record Route: NONE
  Fspec: ave rate=0 kbits, burst=1000 bytes, peak rate=0 kbits
Shortest Unconstrained Path Info:
  Path Weight: 1 (TE)
  Explicit Route: 30.30.30.5 30.30.30.6 6.6.6.6
History:
Tunnel:
  Time since created: 4 hours, 53 minutes
  Time since path change: 4 hours, 52 minutes
  Number of LSP IDs (Tun_Instances) used: 14
Current LSP:
  Uptime: 4 hours, 52 minutes
```

Figure 7 - Checking Backup Tunnel 100 on P1

## 3. Testing MPLS TE Fast Reroute Link Protection

We now disable the Gi0/1 interface on P1 to simulate a link failure between P1 and P3:

```
P1(config)# int gi0/1
P1(config-if)# sh
```

The backup path via P2 is now activated, and the status changes from ready to active (Figure 8).

```
P1#show mpls traffic-eng fast-reroute database
Headend frr information:
Protected tunnel          In-label Out intf/label  FRR intf/label  Status

LSP midpoint frr information:
LSP identifier          In-label Out intf/label  FRR intf/label  Status
2.2.2.2 12 [15]        17      Gi0/1:16      Tu100:16      active
P1#
```

Figure 8 - Checking MPLS TE FRR Database After Link Failure Between P1 and P3

Figure 9 shows an additional label 18, that was inserted into the network traffic by P1.

```
PE1# traceroute 192.168.2.1
```



# MPLS Traffic Engineering

```

PE1#traceroute 192.168.2.1
Type escape sequence to abort.
Tracing the route to 192.168.2.1
VRF info: (vrf in name/id, vrf out name/id)
 1 20.20.20.2 [MPLS: Label 17 Exp 0] 16 msec 3 msec 4 msec
 2 30.30.30.2 [MPLS: Labels 18/16 Exp 0] 12 msec 7 msec 7 msec
 3 40.40.40.10 [MPLS: Label 16 Exp 0] 10 msec 9 msec 15 msec
 4 40.40.40.6 [MPLS: Label 17 Exp 0] 13 msec 7 msec 6 msec
 5 50.50.50.2 14 msec 8 msec *
PE1#
    
```

Figure 9 - Checking Connectivity between PE1 and PE2 After Link Failure on P1

Indeed, a screenshot of the packets captured on the link between P1 and P2 shows that both labels 18 and 16 are presented in the traffic (Figure 10).

```

-Frame 147: 122 bytes on wire (976 bits), 122 bytes captured (976 bits) on interface -, id 0
-Ethernet II, Src: 0c:18:b8:69:00:02 (0c:18:b8:69:00:02), Dst: 0c:14:cb:31:00:00 (0c:14:cb:31:00:00)
-MultiProtocol Label Switching Header, Label: 18, Exp: 0, S: 0, TTL: 254
-MultiProtocol Label Switching Header, Label: 16, Exp: 0, S: 1, TTL: 254
-Internet Protocol Version 4, Src: 2.2.2.2, Dst: 192.168.2.1
-Internet Control Message Protocol
    
```

Figure 10 - Capture Packets between PE1 and PE2 After Link Failure on P1

Notice the primary (red and the backup (green) paths and particular labels assignment depicted in Figure 11.




Figure 11 - MPLS Label Assignment on Backup Path from PE to PE2

At P2, label 18 is popped, and the traffic is subsequently redirected to P3 with only label 16, as depicted in Figure 12.

```

P2#show mpls forwarding-table
Local   Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label   Label     or Tunnel Id   Switched     interface
16      16        8.8.8.8 21 [11] 1180        Gi0/0      30.30.30.1
17      Pop Label  6.6.6.6 100 [11] 0           Gi0/0      30.30.30.1
18      Pop Label  4.4.4.4 100 [14] 9348        Gi0/1      40.40.40.10
    
```

Figure 12 - Popping Label 18 on P2

Finally, the backup LSP is merged with the primary LSP at the Merge Point, which is router P3.

### Conclusion

MPLS Traffic Engineering (TE) provides a way to optimize network traffic and ensure efficient use of network resources.

Using a Merge Point (MP) and Fast Reroute (FRR) techniques in MPLS TE enables network operators to minimize the impact of link and node failures on traffic and ensure faster restoration of services.

The MP is a critical component of FRR, ensuring that traffic from the primary and backup paths is seamlessly merged without any packet loss or duplication.

Network operators need to understand the concept of MP, FRR, and MPLS TE to deploy these technologies effectively and enhance the overall resiliency of their networks. By leveraging these techniques, network operators can ensure that their networks are always available and provide high-quality services to their customers.





This ebook was brought to you by [Noction](#)

---

Copyright © Noction Inc., All Rights Reserved. Noction logos, and trademarks or registered trademarks of Noction Inc. or its subsidiaries in the United States and other countries.

Other names and brands may be claimed as the property of others. Information regarding third party products is provided solely for educational purposes.

Noction Inc. is not responsible for the performance or support of third party products and does not make any representations or warranties whatsoever regarding quality, reliability, functionality, or compatibility of these devices or products.

**Copyright © Noction Inc.**

