

CCIE Service Provider Lab Workbook v4.0

(<http://labs.ine.com/workbook/toc/service-provider-v4>) »

CCIE SP v4 Advanced Technology Labs - MPLS TE

MPLS TE Fast Reroute

« [MPLS TE with Targeted LDP Adjacencies \(/workbook/view/service-provider-v4/task/mpls-te-with-targeted-ldp-adjacencies-Mjg4NQ%3D%3D\)](#) | [Inter-Area MPLS TE with OSPF \(/workbook/view/service-provider-v4/task/inter-area-mpls-te-with-ospf-Mjg4Nw%3D%3D\)](#) »

Last updated: April 22, 2016

Note:

Initial Configuration & Diagrams: [Load the initial configuration files for the section named MPLS TE with IS-IS, which can be found in CCIE SPv4 Topology Diagrams & Initial Configurations \(<http://labs.ine.com/workbook/view/service-provider-v4/task/ccie-spv4-topology-diagrams-initial-configs>\).](#) [Refer to the Base IPv4 Diagram in order to complete this task.](#)

Task

- R2 and XR1 are preconfigured as PE routers for the MPLS L3VPN customer routers R1 and XR2, respectively, but the core of the Service Provider network is not running LDP.
- Configure the core of the Service Provider network to support MPLS TE tunnels as follows:
 - Enable MPLS TE support for the IS-IS Level 2 core.
 - Set the IS-IS MPLS TE Router-ID to be the Loopback0 interfaces.
 - Enable support for RSVP and MPLS TE on all transit interfaces running IS-IS in the core.
- Configure an MPLS TE tunnel from R2 to XR1 as follows:
 - Unnumber the tunnel to R2's Loopback0 interface.
 - Set the tunnel destination as XR1's Loopback0 interface.
 - Set the tunnel's path option to explicitly follow the path from R2 to R3 to R4 to R5 to XR1.
 - Configure Autoroute Announce on the tunnel so that the IS-IS core can use it for dynamic routing.
- Configure an MPLS TE tunnel from XR1 to R2 as follows:
 - Unnumber the tunnel to XR1's Loopback0 interface.
 - Set the tunnel destination as R2's Loopback0 interface.
 - Set the tunnel's path option to explicitly follow the path from XR1 to R5 to R6 to R3 to R2.
 - Configure Autoroute Announce on the tunnel so that the IS-IS core can use it for dynamic routing.
- Configure an MPLS TE tunnel from R4 to R5 for FRR as follows:
 - Unnumber the tunnel to R4's Loopback0 interface.
 - Set the tunnel destination as R5's Loopback0 interface.
 - Set the tunnel's path option to explicitly avoid the link between R4 and R5.
 - Configure BFD signaling between R4 and R5 to detect a link failure between them in less than one second.
 - R2's TE tunnel to XR1 should be Fast Rerouted if the link between R4 and R5 is down.
- Configure an MPLS TE tunnel from R6 to R3 for FRR as follows:

<https://t.me/learningnets>

- Unnumber the tunnel to R6's Loopback0 interface.
- Set the tunnel destination as R3's Loopback0 interface.
- Set the tunnel's path option to explicitly avoid the link between R3 and R6.
- Configure BFD signaling between R3 and R6 to detect a link failure between them in less than one second.
- XR1's TE tunnel to R2 should be Fast Rerouted if the link between R3 and R6 is down.
- When complete, you should be able to perform the following verifications:
 - Traceroutes from R1 to XR2 should follow the R2 to R3 to R4 to R5 to XR1 path.
 - Remove VLAN 45 from Gig1.45 on R5 to simulate a link failure between R4 and R5; immediately following this the traceroutes from R1 to XR2 should be Fast Rerouted via R6.
 - Traceroutes from XR2 to R1 should follow the XR1 to R5 to R6 to R3 to R1 path.
 - Remove VLAN 36 from Gig1.36 on R3 to simulate a link failure between R3 and R6;; immediately following this the traceroutes from XR2 to R1 should be Fast Rerouted via R4.

Configuration [Click to collapse](#)

```

R2:
mpls traffic-eng tunnels
!
ip explicit-path name TO_XR1 enable
  next-address 3.3.3.3
  next-address 4.4.4.4
  next-address 5.5.5.5
  next-address 19.19.19.19
!
interface Tunnel0
  ip unnumbered Loopback0
  tunnel mode mpls traffic-eng
  tunnel destination 19.19.19.19
  tunnel mpls traffic-eng autoroute announce
  tunnel mpls traffic-eng path-option 1 explicit name TO_XR1
  tunnel mpls traffic-eng fast-reroute
!
interface GigabitEthernet1.23
  mpls traffic-eng tunnels
  ip rsvp bandwidth
!
interface GigabitEthernet1.24
  mpls traffic-eng tunnels
  ip rsvp bandwidth
!

```

```

router isis
  metric-style wide
  mpls traffic-eng router-id Loopback0
  mpls traffic-eng level-2

R3:
mpls traffic-eng tunnels
!
interface GigabitEthernet1.23
  mpls traffic-eng tunnels
  ip rsvp bandwidth
!
interface GigabitEthernet1.34
  mpls traffic-eng tunnels
  ip rsvp bandwidth
!
interface GigabitEthernet1.36
  mpls traffic-eng tunnels
  bfd interval 250 min_rx 250 multiplier 3
  ip rsvp bandwidth
  ip rsvp signalling hello bfd
!
router isis
  metric-style wide
  mpls traffic-eng router-id Loopback0
  mpls traffic-eng level-2
!
ip rsvp signalling hello bfd

```

```

R4:
mpls traffic-eng tunnels
!
interface Tunnel1
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 5.5.5.5
 tunnel mpls traffic-eng path-option 1 explicit name AVOID_R4_R5_LINK
!
interface GigabitEthernet1.24
 mpls traffic-eng tunnels
 ip rsvp bandwidth
!
interface GigabitEthernet1.34
 mpls traffic-eng tunnels
 ip rsvp bandwidth
!
interface GigabitEthernet1.45
 mpls traffic-eng tunnels
 mpls traffic-eng backup-path Tunnel1
 bfd interval 250 min_rx 250 multiplier 3
 ip rsvp bandwidth
 ip rsvp signalling hello bfd
!
interface GigabitEthernet1.46

```

```

 mpls traffic-eng tunnels
 ip rsvp bandwidth
!
router isis
 metric-style wide
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-2
!
 ip rsvp signalling hello bfd
!
 ip explicit-path name AVOID_R4_R5_LINK enable
 exclude-address 20.4.5.4
 exclude-address 20.4.5.5

```

```

R5:
mpls traffic-eng tunnels
!
interface GigabitEthernet1.45
 mpls traffic-eng tunnels
 bfd interval 250 min_rx 250 multiplier 3
 ip rsvp bandwidth
 ip rsvp signalling hello bfd
!
interface GigabitEthernet1.56
 mpls traffic-eng tunnels
 ip rsvp bandwidth
!
interface GigabitEthernet1.519

```

```

mpls traffic-eng tunnels

ip rsvp bandwidth

!

router isis

metric-style wide

mpls traffic-eng router-id Loopback0

mpls traffic-eng level-2

!

ip rsvp signalling hello bfd

R6:

mpls traffic-eng tunnels

!

interface Tunnel1

ip unnumbered Loopback0

tunnel mode mpls traffic-eng

tunnel destination 3.3.3.3

tunnel mpls traffic-eng path-option 1 explicit name AVOID_R3_R6_LINK

!

interface GigabitEthernet1.36

mpls traffic-eng tunnels

ip rsvp bandwidth

mpls traffic-eng backup-path Tunnel1

bfd interval 250 min_rx 250 multiplier 3

ip rsvp signalling hello bfd

!

```

```

interface GigabitEthernet1.46

mpls traffic-eng tunnels

ip rsvp bandwidth

!

interface GigabitEthernet1.56

mpls traffic-eng tunnels

ip rsvp bandwidth

!

interface GigabitEthernet1.619

mpls traffic-eng tunnels

ip rsvp bandwidth

!

router isis

metric-style wide

mpls traffic-eng router-id Loopback0

mpls traffic-eng level-2

!

ip rsvp signalling hello bfd

!

ip explicit-path name AVOID_R3_R6_LINK enable

exclude-address 20.3.6.6

exclude-address 20.3.6.3

XR1:

explicit-path name TO_R2

index 1 next-address strict ipv4 unicast 5.5.5.5

index 2 next-address strict ipv4 unicast 6.6.6.6

index 3 next-address strict ipv4 unicast 3.3.3.3

```

```

index 4 next-address strict ipv4 unicast 2.2.2.2
!
interface tunnel-te0
  ipv4 unnumbered Loopback0
  autoroute announce
  destination 2.2.2.2
  path-option 1 explicit name TO_R2
  fast-reroute
!
router isis 1
  address-family ipv4 unicast
  metric-style wide
  mpls traffic-eng level-2-only
  mpls traffic-eng router-id Loopback0
!
rsvp
  interface GigabitEthernet0/0/0/0.519
  !
  interface GigabitEthernet0/0/0/0.619
  !
mpls traffic-eng
  interface GigabitEthernet0/0/0/0.519
  !
  interface GigabitEthernet0/0/0/0.619
  !
mpls ldp

```

Verification

Prior to any failures in the network, traffic from R1 to XR2 follows the explicit route of the TE tunnel from R2 to XR1.

```

R1#traceroute 20.20.20.20

Type escape sequence to abort.
Tracing the route to 20.20.20.20
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.2.2 4 msec 1 msec 1 msec
 2 20.2.3.3 [MPLS: Labels 29/16018 Exp 0] 32 msec 25 msec 16 msec
 3 20.3.4.4 [MPLS: Labels 19/16018 Exp 0] 26 msec 20 msec 32 msec
 4 20.4.5.5 [MPLS: Labels 18/16018 Exp 0] 30 msec 130 msec 49 msec
 5 20.5.19.19 21 msec 23 msec 136 msec
 6 10.19.20.20 57 msec * *

```

R2 requests FRR protection for this tunnel.

```
R2#show mpls traffic-eng tunnels protection
```

```
P2P TUNNELS:
```

```
R2_t0
```

```
LSP Head, Tunnel0, Admin: up, Oper: up
```

```
Src 2.2.2.2, Dest 19.19.19.19, Instance 28
```

```
Fast Reroute Protection: Requested
```

```
Outbound: Unprotected: no backup tunnel assigned
```

```
LSP signalling info:
```

```
Original: out i/f: Gi1.23, label: 29, nhop: 20.2.3.3
```

```
nnhop: 4.4.4.4; nnhop rtr id: 4.4.4.4
```

```
Path Protection: None
```

```
P2MP TUNNELS:
```

R4 and R5 are BFD adjacent. Note that the “client” protocol of BFD is Fast Reroute (FRR). Also notice that the R2-XR1 tunnel is currently being protected by Tunnel1 on R4.

```
R4#show ip rsvp hello bfd nbr detail
```

```
Hello Client Neighbors
```

```
Remote addr 20.4.5.5, Local addr 20.4.5.4
```

```
Type: Active
```

```
I/F: Gi1.45
```

```
State: Up (for 00:00:25)
```

```
Clients: FRR
```

```
LSPs protecting: 1 (frr: 1, hst upstream: 0 hst downstream: 0)
```

```
Communication with neighbor lost: 3
```

```
R4#show mpls traffic-eng fast-reroute database
```

```
P2P Headend FRR information:
```

Protected tunnel	In-label	Out intf/label	FRR intf/label	Status
-----	-----	-----	-----	-----

```
P2P LSP midpoint frr information:
```

LSP identifier	In-label	Out intf/label	FRR intf/label	Status
-----	-----	-----	-----	-----
2.2.2.2 0 [32]	17	Gi1.45:17	Tu1:17	ready

```
P2MP Sub-LSP FRR information:
```

```
*Sub-LSP identifier
```

src_lspid[subid]->dst_tunid	In-label	Out intf/label	FRR intf/label	Status
-----	-----	-----	-----	-----

```
* Sub-LSP identifier format: <TunSrc>_<LSP_ID>[SubgroupID]-><TunDst>_<Tun_ID>
```

```
Note: Sub-LSP identifier may be truncated.
```

```
Use 'detail' display for the complete key.
```

R5 removes VLAN 45 from Gig1.45, which causes a loss of communication on the link between R4 and R5.

```
R5(config-subif)#interface GigabitEthernet1.45
R5(config-subif)#no encapsulation dot1q 45
% IP address 20.4.5.5 will be removed from GigabitEthernet1.45 due to removal of VLAN encapsulation.
```

R4 detects that the BFD neighbor R5 is down, which triggers the FRR protection.

```
R4#show bfd neighbor

IPv4 Sessions

NeighAddr                LD/RD          RH/RS    State    Int
20.4.5.5                  4097/0        Down     Down     Gi1.45
```

The traceroute from R1 to XR2 now indicates that the path is being rerouted from R4 to R5 to R6. Note the additional label in the stack, which is required to tunnel the R4 to R5 traffic over R6.

```
R1#traceroute 20.20.20.20
Type escape sequence to abort.
Tracing the route to 20.20.20.20
VRF info: (vrf in name/id, vrf out name/id)

 1 10.1.2.2 4 msec 1 msec 6 msec
 2 *
   20.2.3.3 [MPLS: Labels 33/16018 Exp 0] 9 msec 8 msec
 3 20.3.4.4 [MPLS: Labels 17/16018 Exp 0] 20 msec 32 msec 36 msec
 4 20.4.6.6 [MPLS: Labels 27/20/16018 Exp 0] 32 msec 37 msec 33 msec
 5 20.5.6.5 [MPLS: Labels 20/16018 Exp 0] 21 msec 21 msec 45 msec
 6 20.5.19.19 26 msec 14 msec 15 msec
 7 10.19.20.20 15 msec * 10 msec
```

Note that R4 is actively protecting R2's tunnel. The previous output of this command showed "Ready", now it shows as "Active", meaning R4 performed local protection and is rerouting R2's tunnel using its local Tunnel1.

```
R4#show mpls traffic-eng fast-reroute database

P2P Headend FRR information:
Protected tunnel          In-label Out intf/label  FRR intf/label  Status
-----
P2P LSP midpoint frr information:
LSP identifier           In-label Out intf/label  FRR intf/label  Status
-----
2.2.2.2 0 [54]          17      Gi1.45:20      Tu1:20          active

P2MP Sub-LSP FRR information:
*Sub-LSP identifier
src_lspid[subid]->dst_tunid  In-label Out intf/label  FRR intf/label  Status
-----

* Sub-LSP identifier format: <TunSrc>_<LSP_ID>[SubgroupID]-><TunDst>_<Tun_ID>
Note: Sub-LSP identifier may be truncated.
Use 'detail' display for the complete key.
```

In the reverse direction, the same occurs. XR2's traceroute to R1 follows XR1's TE tunnel's explicit path.

```
RP/0/0/CPU0:XR2#traceroute 1.1.1.1
```

```
Sun May 31 14:14:34.313 UTC
```

```
Type escape sequence to abort.
```

```
Tracing the route to 1.1.1.1
```

```
 1 10.19.20.19 0 msec  0 msec  0 msec
 2 20.5.19.5 [MPLS: Labels 19/17 Exp 0] 0 msec  0 msec  0 msec
 3 20.5.6.6 [MPLS: Labels 30/17 Exp 0] 0 msec  89 msec  39 msec
 4 20.3.6.3 [MPLS: Labels 30/17 Exp 0] 19 msec  19 msec  19 msec
 5 10.1.2.2 [MPLS: Label 17 Exp 0] 19 msec  19 msec  19 msec
 6 10.1.2.1 19 msec  *  9 msec
```

XR1 has requested that this tunnel be FRR protected.

```
RP/0/0/CPU0:XR1#show mpls traffic-eng tunnels protection
```

```
Sun May 31 14:14:55.262 UTC
```

```
R2_t0 Tunnel Id: 0
```

```
LSP Tail, signaled, connection up
```

```
Src: 2.2.2.2, Dest: 19.19.19.19, Instance: 54
```

```
Fast Reroute Protection: Requested
```

```
Inbound: FRR Inactive
```

```
LSP signalling info:
```

```
Original: in i/f: GigabitEthernet0/0/0.519, label: 3, phop: 20.5.19.5
```

```
XR1_t0 Tunnel Id: 0
```

```
LSP Head, Admin: up, Oper: up
```

```
Src: 19.19.19.19, Dest: 2.2.2.2, Instance: 34
```

```
Fast Reroute Protection: Requested
```

```
Outbound: FRR Inactive
```

```
LSP signalling info:
```

```
Original: out i/f: GigabitEthernet0/0/0.519, label: 19, nhop: 20.5.19.5
```

R3 and R6 are BFD adjacent, and R6 has registered this adjacency with the FRR process.

```
R6#show ip rsvp hello bfd nbr detail
```

```
Hello Client Neighbors
```

```
Remote addr 20.3.6.3, Local addr 20.3.6.6
```

```
Type: Active
```

```
I/F: Gi1.36
```

```
State: Up (for 00:04:43)
```

```
Clients: FRR
```

```
LSPs protecting: 1 (frr: 1, hst upstream: 0 hst downstream: 0)
```

```
Communication with neighbor lost: 65
```

R3 removes VLAN36 from Gig1.36, which causes communication to be lost between R3 and R6.

```
R3(config)#interface GigabitEthernet1.36
```

```
R3(config-subif)#no encapsulation dot1q 36
```

```
% IP address 20.3.6.3 will be removed from GigabitEthernet1.36 due to removal of VLAN encapsulation.
```

When R6 loses the BFD neighbor R3, the FRR process is triggered.

<https://t.me/learningnets>

```

R6#show ip rsvp hello bfd nbr detail

Hello Client Neighbors

Remote addr 20.3.6.3, Local addr 20.3.6.6

Type: Passive

I/F: Gi1.36

State: Lost (for 00:00:24)

Clients: None

LSPs protecting: 0

Communication with neighbor lost: 66

```

The traceroute from XR2 to R1 now indicates that the tunnel is being rerouted via R4. Again, note the additional label in the stack that is needed for the FRR tunnel between R6 and R3.

```

RP/0/0/CPU0:XR2#traceroute 1.1.1.1

Sun May 31 14:16:10.067 UTC

Type escape sequence to abort.

Tracing the route to 1.1.1.1

 1 10.19.20.19 0 msec  0 msec  0 msec
 2  *
   20.5.19.5 9 msec  0 msec
 3 20.5.6.6 [MPLS: Labels 30/17 Exp 0] 0 msec  0 msec  0 msec
 4 20.4.6.4 [MPLS: Labels 16/30/17 Exp 0] 0 msec  0 msec  0 msec
 5 20.3.4.3 [MPLS: Labels 30/17 Exp 0] 0 msec  0 msec  0 msec
 6 10.1.2.2 [MPLS: Label 17 Exp 0] 0 msec  0 msec  0 msec
 7 10.1.2.1 9 msec  *  9 msec

```

In most fast-reroute designs, midpoint routers such as R4 and R6 are the devices that provide fast reroute local protection for tunnels originating at PE routers, such as in this example. As soon as a failure is detected, a midpoint router is able to quickly repair the tunnel, giving the headend router enough time to re-compute SPF and find an alternative best path. Note that in our design, R2 and XR1 will never try to find an alternate path, as they only have a single explicit path option defined. If we configured a less preferred dynamic option on these tunnels, R2 and XR1 would then re-compute CSPF and find an alternate best path. However, for the meantime while they find an alternate path, the fast-reroute tunnel is preventing packet loss by re-routing their active tunnels around a failure.

There are generally two types of failures that can be protected with fast-reroute tunnels; link and node failures. This example showed how to route around a link failure, by configuring the fast-reroute tunnel with an explicit path that completely avoids the link that is being protected (Gig1.36 and Gig1.45). Re-routing around a link does not help if an entire node goes down, such as if R5 crashes and reboots. R4 could have configured its re-route tunnel to complete avoid R5 and go to XR1 directly, and this would have been considered "node" protection.

Another benefit of using local protection on midpoint routers is that they can protect multiple tunnels at the same time. If R2 had several tunnels to XR1, or if there were several PEs with tunnels crossing the Gig1.45 link, R4 would be able to protect all of those tunnels with the current configuration. This is called 1:N protection, where a single fast-reroute tunnel can protect several (N) primary tunnels.

The mechanics behind fast-reroute involve label stacking, as can be seen from the traceroute outputs. Lets revert the broken links to a working configuration and follow the labels.

```

R3#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R3(config)#interface GigabitEthernet1.36
R3(config-subif)# encapsulation dot1Q 36
R3(config-subif)# ip address 20.3.6.3 255.255.255.0
R3(config-subif)# ip router isis
R3(config-subif)# mpls traffic-eng tunnels
R3(config-subif)# bfd interval 250 min_rx 250 multiplier 3
R3(config-subif)# ip rsvp bandwidth
R3(config-subif)# ip rsvp signalling hello bfd
R3(config-subif)#end
R3#
%SYS-5-CONFIG_I: Configured from console by console

```

```

R5#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R5(config)#
%SYS-5-CONFIG_I: Configured from console by console
R5(config)#interface GigabitEthernet1.45
R5(config-subif)# encapsulation dot1Q 45
R5(config-subif)# ip address 20.4.5.5 255.255.255.0
R5(config-subif)# ip router isis
R5(config-subif)# mpls traffic-eng tunnels
R5(config-subif)# bfd interval 250 min_rx 250 multiplier 3
R5(config-subif)# ip rsvp bandwidth
R5(config-subif)# ip rsvp signalling hello bfd
R5(config-subif)#end
R5#
%SYS-5-CONFIG_I: Configured from console by console

```

Under normal circumstances, without a failure in the network, the LSP from R2 to XR1 uses the following labels:

```

R2#traceroute 19.19.19.19
Type escape sequence to abort.
Tracing the route to 19.19.19.19
VRF info: (vrf in name/id, vrf out name/id)
 1 20.2.3.3 [MPLS: Label 32 Exp 0] 12 msec 7 msec 5 msec
 2 20.3.4.4 [MPLS: Label 18 Exp 0] 21 msec 30 msec 35 msec
 3 20.4.5.5 [MPLS: Label 21 Exp 0] 30 msec 34 msec 30 msec
 4 20.5.19.19 22 msec * 10 msec

```

R2 uses label 32 to reach R3, R3 uses label 18 to reach R4, R4 uses label 21 to reach R5, and R5 pops the label towards XR1.

This can be verified by looking at the LFIB:

```

R2#show mpls forwarding-table 19.19.19.19
Local      Outgoing Prefix          Bytes Label  Outgoing  Next Hop
Label      Label      or Tunnel Id    Switched     interface
None [T]  32         19.19.19.19/32  0            Tu0        point2point

[T] Forwarding through a LSP tunnel.
View additional labelling info with the 'detail' option

R3#show mpls forwarding-table labels 32
Local      Outgoing Prefix          Bytes Label  Outgoing  Next Hop
Label      Label      or Tunnel Id    Switched     interface
32         18         2.2.2.2 0 [119] 1238        Gi1.34     20.3.4.4

R4#show mpls forwarding-table label 18
Local      Outgoing Prefix          Bytes Label  Outgoing  Next Hop
Label      Label      or Tunnel Id    Switched     interface
18         21         2.2.2.2 0 [119] 1678        Gi1.45     20.4.5.5

R5#show mpls forwarding-table labels 21
Local      Outgoing Prefix          Bytes Label  Outgoing  Next Hop
Label      Label      or Tunnel Id    Switched     interface
21         Pop Label  2.2.2.2 0 [119] 1942        Gi1.519    20.5.19.19

```

In order for R4 to provide local repair of tunnels using its Gig1.45, it must encapsulate packets in a way such that R5 receives a label stack with label 21 as the topmost label. This way R5 receives labeled packets as if its Gig1.45 had never failed, and is then able to switch the packets towards XR1 by popping the label, as its current LFIB is programmed to do.

R4's fast-reroute tunnel avoids Gig1.45 by going through R6 first (Gig1.45), then to R5 (Gig1.56). Label 27 is pushed as packets are encapsulate into this tunnel.

```
R4#show mpls traffic-eng tunnels
```

```
P2P TUNNELS/LSPs:
```

```
Name: R4_t1 (Tunnel) Destination: 5.5.5.5
Status:
Admin: up Oper: up Path: valid Signalling: connected
path option 1, type explicit AVOID_R4_R5_LINK (Basis for Setup, path weight 20)
```

```
Config Parameters:
```

```
Bandwidth: 0 kbps (Global) Priority: 7 7 Affinity: 0x0/0xFFFF
Metric Type: TE (default)
AutoRoute: disabled LockDown: disabled Loadshare: 0 [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 : GigabitEthernet1.46, 27
```

```
Next Hop : 20.4.6.6
```

```
RSVP Signalling Info:
```

```
Src 4.4.4.4, Dst 5.5.5.5, Tun_Id 1, Tun_Instance 7
```

```
RSVP Path Info:
```

```
My Address: 20.4.6.4
```

```
Explicit Route: 20.4.6.6 20.5.6.6 20.5.6.5 5.5.5.5
```

```
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
```

```
History:
```

```
Tunnel:
```

```
Time since created: 52 minutes, 45 seconds
```

```
Time since path change: 51 minutes, 57 seconds
```

```
Number of LSP IDs (Tun_Instances) used: 7
```

```
Current LSP: [ID: 7]
```

```
Uptime: 51 minutes, 57 seconds
```

When R6 receives packets labeled with 27, it pops the label and forwards them to R5:

```
R6#show mpls forwarding-table labels 27
```

Local Label	Outgoing Label	Prefix or Tunnel Id	Bytes Switched	Outgoing interface	Next Hop
27	Pop Label	4.4.4.4 1 [7]	37524	Gi1.56	20.5.6.5

This means that when Gig1.45 fails, R4 will push label 21 as if it were encapsulating packets out of Gig1.45, and then pushes label 27 on top of this label. This is the "fast-reroute" label. R6 receives packets with label 27, pops it, and sends the labeled packet to R5. R5 receives the packet with label 21 as if it had received them from Gig1.45, pops the label, and sends the packet to XR1. This clever use of label stacking is the mechanism behind fast-reroute.

```
R5(config)#interface GigabitEthernet1.45
R5(config-subif)#no encapsulation dot1Q 45

% IP address 20.4.5.5 will be removed from GigabitEthernet1.45 due to removal of VLAN encapsulation.

R2#traceroute 19.19.19.19

Type escape sequence to abort.
Tracing the route to 19.19.19.19
VRF info: (vrf in name/id, vrf out name/id)
 1 20.2.3.3 [MPLS: Label 32 Exp 0] 8 msec 6 msec 7 msec
 2 20.3.4.4 [MPLS: Label 18 Exp 0] 24 msec 30 msec 47 msec
 3 20.4.6.6 [MPLS: Labels 27/21 Exp 0] 19 msec 31 msec 31 msec
 4 20.5.6.5 [MPLS: Label 21 Exp 0] 23 msec 30 msec 30 msec
 5 20.5.19.19 19 msec * 9 msec
```

There is an alternate way of configuring backup tunnels called AutoTunnel Backup. Next-hop (link protection) and Next-Next-Hop (node protection) tunnels are configured automatically on each router. If this feature were enabled on R3, R4, R5, and R6, all links and nodes would be protected in the SP core.

Remove the manually configured tunnels from R4 and R6.

```
R4(config)#no interface Tunnel1

R6(config)#no interface Tunnel1
```

Note that all pre-requisites for establishing a TE tunnel must be previously enabled, such as enabling TE extensions for IGP and configuring TE on the interfaces.

```
R3-R6:

mpls traffic-eng auto-tunnel backup

mpls traffic-eng auto-tunnel backup config unnumbered-interface Loopback 0
```

After enabling Auto-Tunnel backup, the core routers we can see that several tunnels have been automatically created:

R3#show mpls traffic-eng tunnels brief

Signalling Summary:

LSP Tunnels Process: running
 Passive LSP Listener: running
 RSVP Process: running
 Forwarding: enabled

auto-tunnel:

backup Enabled (3), id-range:65436-65535
 onehop Disabled (0), id-range:65336-65435
 mesh Disabled (0), id-range:64336-65335

Periodic reoptimization: every 3600 seconds, next in 1250 seconds

Periodic FRR Promotion: Not Running

Periodic auto-tunnel:

backup notinuse scan: every 3600 seconds, next in 2456 seconds
 Periodic auto-bw collection: every 300 seconds, next in 50 seconds

P2P TUNNELS/LSPs:

TUNNEL NAME	DESTINATION	UP IF	DOWN IF	STATE/PROT
R3_t65436	2.2.2.2	-	Gi1.34	up/up
R3_t65437	4.4.4.4	-	Gi1.23	up/up
R3_t65438	5.5.5.5	-	Gi1.36	up/up
R2_t0	19.19.19.19	Gi1.23	Gi1.34	up/up
R5_t65437	3.3.3.3	Gi1.34	-	up/up
R6_t65436	3.3.3.3	Gi1.34	-	up/up
XR1_t0	2.2.2.2	Gi1.36	Gi1.23	up/up

Displayed 3 (of 3) heads, 2 (of 2) midpoints, 2 (of 2) tails

P2MP TUNNELS:

Displayed 0 (of 0) P2MP heads

P2MP SUB-LSPS:

Displayed 0 P2MP sub-LSPs:

0 (of 0) heads, 0 (of 0) midpoints, 0 (of 0) tails

R4#show mpls traffic-eng tunnels brief

Signalling Summary:

LSP Tunnels Process: running
 Passive LSP Listener: running
 RSVP Process: running
 Forwarding: enabled

auto-tunnel:

backup Enabled (2), id-range:65436-65535
 onehop Disabled (0), id-range:65336-65435
 mesh Disabled (0), id-range:64336-65335

Periodic reoptimization: every 3600 seconds, next in 1273 seconds

Periodic FRR Promotion: Not Running

Periodic auto-tunnel:

backup notinuse scan: every 3600 seconds, next in 2420 seconds
 Periodic auto-bw collection: every 300 seconds, next in 73 seconds

P2P TUNNELS/LSPs:

TUNNEL NAME	DESTINATION	UP IF	DOWN IF	STATE/PROT
R4_t65436	5.5.5.5	-	Gi1.46	up/up
R4_t65437	19.19.19.19	-	Gi1.46	up/up
R2_t0	19.19.19.19	Gi1.34	Gi1.45	up/up
R3_t65436	2.2.2.2	Gi1.34	Gi1.24	up/up
R3_t65437	4.4.4.4	Gi1.24	-	up/up
R5_t65436	6.6.6.6	Gi1.45	Gi1.46	up/up
R5_t65437	3.3.3.3	Gi1.45	Gi1.34	up/up
R6_t65436	3.3.3.3	Gi1.46	Gi1.34	up/up
R6_t65437	2.2.2.2	Gi1.46	Gi1.24	up/up

Displayed 2 (of 2) heads, 6 (of 6) midpoints, 1 (of 1) tails

P2MP TUNNELS:

Displayed 0 (of 0) P2MP heads

P2MP SUB-LSPS:

Displayed 0 P2MP sub-LSPs:

0 (of 0) heads, 0 (of 0) midpoints, 0 (of 0) tails

R5#show mpls traffic-eng tunnels brief

Signalling Summary:

LSP Tunnels Process: running
 Passive LSP Listener: running
 RSVP Process: running

Forwarding: enabled

auto-tunnel:

backup Enabled (3), id-range:65436-65535
 onehop Disabled (0), id-range:65336-65435
 mesh Disabled (0), id-range:64336-65335

Periodic reoptimization: every 3600 seconds, next in 1271 seconds

Periodic FRR Promotion: Not Running

Periodic auto-tunnel:

backup notinuse scan: every 3600 seconds, next in 2410 seconds

Periodic auto-bw collection: every 300 seconds, next in 71 seconds

P2P TUNNELS/LSPs:

TUNNEL NAME	DESTINATION	UP IF	DOWN IF	STATE/PROT
R5_t65436	6.6.6.6	-	Gi1.45	up/up
R5_t65437	3.3.3.3	-	Gi1.45	up/up
R5_t65438	19.19.19.19	-	Gi1.56	up/up
R2_t0	19.19.19.19	Gi1.45	Gi1.519	up/up
R3_t65438	5.5.5.5	Gi1.56	-	up/up
R4_t65436	5.5.5.5	Gi1.56	-	up/up
XR1_t0	2.2.2.2	Gi1.519	Gi1.56	up/up

Displayed 3 (of 3) heads, 2 (of 2) midpoints, 2 (of 2) tails

P2MP TUNNELS:

Displayed 0 (of 0) P2MP heads

P2MP SUB-LSPS:

Displayed 0 P2MP sub-LSPs:

0 (of 0) heads, 0 (of 0) midpoints, 0 (of 0) tails

```
R6#show mpls traffic-eng tunnels brief
```

Signalling Summary:

```
LSP Tunnels Process:      running
Passive LSP Listener:    running
RSVP Process:            running
Forwarding:              enabled
```

auto-tunnel:

```
  backup Enabled (2 ), id-range:65436-65535
  onehop Disabled (0 ), id-range:65336-65435
  mesh   Disabled (0 ), id-range:64336-65335
```

```
Periodic reoptimization:  every 3600 seconds, next in 1297 seconds
```

```
Periodic FRR Promotion:   Not Running
```

Periodic auto-tunnel:

```
  backup notinuse scan:   every 3600 seconds, next in 2402 seconds
Periodic auto-bw collection: every 300 seconds, next in 97 seconds
```

P2P TUNNELS/LSPs:

TUNNEL NAME	DESTINATION	UP IF	DOWN IF	STATE/PROT
R6_t65436	3.3.3.3	-	Gi1.46	up/up
R6_t65437	2.2.2.2	-	Gi1.46	up/up
R3_t65438	5.5.5.5	Gi1.36	Gi1.56	up/up
R4_t65436	5.5.5.5	Gi1.46	Gi1.56	up/up
R4_t65437	19.19.19.19	Gi1.46	Gi1.619	up/up
R5_t65436	6.6.6.6	Gi1.46	-	up/up
R5_t65438	19.19.19.19	Gi1.56	Gi1.619	up/up
XR1_t0	2.2.2.2	Gi1.56	Gi1.36	up/up

Displayed 2 (of 2) heads, 5 (of 5) midpoints, 1 (of 1) tails

P2MP TUNNELS:

Displayed 0 (of 0) P2MP heads

P2MP SUB-LSPS:

Displayed 0 P2MP sub-LSPs:

0 (of 0) heads, 0 (of 0) midpoints, 0 (of 0) tails

CONTENTS >

With this configuration, the R2 and XR1 tunnels are protected against any link or node failure in the core network.

For example, R3 is currently protecting R2's tunnel to XR1 which is using Gig1.34, and XR1's tunnel to R2, which is using Gig1.23.

```

R3#show mpls traffic-eng fast-reroute database

P2P Headend FRR information:
Protected tunnel          In-label Out intf/label  FRR intf/label  Status
-----
-----

P2P LSP midpoint frr information:
LSP identifier           In-label Out intf/label  FRR intf/label  Status
-----
-----
2.2.2.2 0 [169]          28      Gi1.34:19      Tu65438:22      ready
19.19.19.19 0 [35]       31      Gi1.23:implicit- Tu65436:implicit ready

P2MP Sub-LSP FRR information:
*Sub-LSP identifier
src_lspid[subid]->dst_tunid  In-label Out intf/label  FRR intf/label  Status
-----
-----

* Sub-LSP identifier format: <TunSrc>_<LSP_ID>[SubgroupID]-><TunDst>_<Tun_ID>
Note: Sub-LSP identifier may be truncated.
Use 'detail' display for the complete key.

```

Tu65438 is providing Node protection for R2's tunnel, by completely avoiding over R4.

```

R3#show derived-config interface Tu65438

Building configuration...

Derived configuration : 217 bytes

!
interface Tunnel65438
 ip unnumbered Loopback0
 no logging event link-status
 tunnel mode mpls traffic-eng
 tunnel destination 5.5.5.5
 tunnel mpls traffic-eng path-option 1 explicit name _dynamic_tunnel65438
end

R3#show ip explicit-path name _dynamic_tunnel65438

PATH _dynamic_tunnel65438 (strict source route, path complete, generation 47, status non-configured)

 1: exclude-address 4.4.4.4

```

Tu65436 is providing link protection for XR1's tunnel by avoiding link Gig1.23.

```
R3#show derived-config interface Tu65436
```

```
Building configuration...
```

```
Derived configuration : 217 bytes
```

```
!  
interface Tunnel65436  
 ip unnumbered Loopback0  
 no logging event link-status  
 tunnel mode mpls traffic-eng  
 tunnel destination 2.2.2.2  
 tunnel mpls traffic-eng path-option 1 explicit name _dynamic_tunnel65436  
end
```

```
R3#show ip explicit-path name _dynamic_tunnel65436
```

```
PATH _dynamic_tunnel65436 (strict source route, path complete, generation 41, status non-configured)
```

```
 1: exclude-address 20.2.3.3
```

« MPLS TE with Targeted LDP Adjacencies (/workbook/view/service-provider-v4/task/mpls-te-with-targeted-ldp-adjacencies-Mjg4NQ%3D%3D) | Inter-Area MPLS TE with OSPF (/workbook/view/service-provider-v4/task/inter-area-mpls-te-with-ospf-Mjg4Nw%3D%3D) »