MPLS TE Bandwidth Reservations

« MPLS TE Explicit Paths (/workbook/view/service-provider-v4/task/mpls-te-explicit-paths-Mjg4MQ%3D%3D) | MPLS TE Metric Manipulation (/workbook/view/service-provider-v4/task/mpls-te-metric-manipulation-Mjg4Mw%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.

CONTENTS

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 R3 and R6 so that 30 Mbps can be reserved via RSVP on any of their interfaces.
  - Configure all other routers in the core so that 750 Mbps can be reserved via RSVP on any of their interfaces.
- 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 to request a bandwidth reservation of 50 Mbps.
  - Set the tunnel's path option to dynamic.
  - 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 to request a bandwidth reservation of 35 Mbps
  - Set the tunnel's path option to dynamic.
  - Configure Autoroute Announce on the tunnel so that the IS-IS core can use it for dynamic routing.
- When complete, the following reachability should be achieved:
  - R1 and XR2 should have full IP reachability to each other, and a traceroute should indicate that their L3VPN tunnel is transiting over the MPLS TE tunnels in the core of the SP network.

Configuration

Click to *collapse*

CONTENTS

```
R2:
mpls traffic-eng tunnels
!
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 bandwidth 50000
 tunnel mpls traffic-eng path-option 1 dynamic
!
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 30000
!
interface GigabitEthernet1.34
 mpls traffic-eng tunnels
 ip rsvp bandwidth 30000
!
interface GigabitEthernet1.36
 mpls traffic-eng tunnels
 ip rsvp bandwidth 30000
!
router isis
 metric-style wide
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-2

R4:
mpls traffic-eng tunnels
!
interface GigabitEthernet1.24
 mpls traffic-eng tunnels
 ip rsvp bandwidth
!
interface GigabitEthernet1.34
 mpls traffic-eng tunnels
```

CONTENTS

```
 ip rsvp bandwidth
!
interface GigabitEthernet1.45
 mpls traffic-eng tunnels
 ip rsvp bandwidth
!
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

R5:
mpls traffic-eng tunnels
!
interface GigabitEthernet1.45
 mpls traffic-eng tunnels
 ip rsvp bandwidth
!
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

R6:
mpls traffic-eng tunnels
!
interface GigabitEthernet1.36
 mpls traffic-eng tunnels
ip rsvp bandwidth 30000
!
interface GigabitEthernet1.46
 mpls traffic-eng tunnels
ip rsvp bandwidth 30000
!
interface GigabitEthernet1.56
 mpls traffic-eng tunnels
ip rsvp bandwidth 30000
!
interface GigabitEthernet1.619
 mpls traffic-eng tunnels
ip rsvp bandwidth 30000
!
```

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

XR1:
interface tunnel-te0
 ipv4 unnumbered Loopback0
 autoroute announce
 signalled-bandwidth 35000
 destination 2.2.2.2
 path-option 1 dynamic
!
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
   bandwidth 750000
 !
 interface GigabitEthernet0/0/0/0.619
!
mpls traffic-eng
 interface GigabitEthernet0/0/0/0.519
 !
 interface GigabitEthernet0/0/0/0.619
 !
mpls ldp
```

CONTENTS ›

Verification

MPLS TE tunnel bandwidth reservations are used to help ensure that a Service Provider can conform to the Service Level Agreements (SLAs) they are selling to customers. Instead of simply using the addition of the bandwidth-based link costs as the deciding path selection metric, like OSPF does, MPLS TE tunnels can use the total bandwidth and available bandwidth values on a per-link basis that RSVP reports as part of the TE topology to calculate the path for an individual tunnel.

The basic logic behind this is that if the Service Provider has 10 customers, each of which has two sites and SLAs that guarantee 1 Gbps of transit bandwidth, the SP can provision 20 MPLS TE tunnels, each requesting 1 Gbps of bandwidth. Remember that because TE tunnels are unidirectional, there must be a tunnel from PE1 to PE2 and from PE2 back to PE1 to provide bidirectional tunneling. When the tunnels are provisioned, the actual path that they take through the SP core can be automatically determined based on the amount of available bandwidth. If the core consists of GigE, OC-48, TenGigE, and HundredGigE links, the tunnels will automatically be arranged in a way that gives each customer a dedicated 1Gbps reservation. The bandwidth available on each link is reported in the Traffic Engineering Database (TED), and the computing node uses this when computing CSPF.

Note, however, that the MPLS TE tunnel bandwidth reservation is only a reservation in the control plane; it is not a reservation in the data plane. This means that there is nothing preventing traffic going over a tunnel from actually using more bandwidth than the reservation says it is guaranteed. You must still enforce at the Service Provider edge the amount of bandwidth that is admitted to the network with some form of admission control, such as policing.

In this example, the core of the network is staged to support reservations up to 750 Mbps everywhere, and reservations up to 30 Mbps on R3 and R6. These values can be verified as follows:

```
R3#show ip rsvp interface
interface    rsvp       allocated  i/f max  flow max sub max  VRF
Gi1          ena        0          750M     750M     0
Gi1.23       ena        0          30M      30M      0
Gi1.34       ena        0          30M      30M      0
Gi1.36       ena        0          30M      30M      0

R4#show ip rsvp interface
interface    rsvp       allocated  i/f max  flow max sub max  VRF
Gi1          ena        0          750M     750M     0
Gi1.24       ena        35M        750M     750M     0
Gi1.34       ena        0          750M     750M     0
Gi1.45       ena        50M        750M     750M     0
Gi1.46       ena        0          750M     750M     0

RP/0/0/CPU0:XR1#show rsvp interface
Fri May 29 00:10:39.110 UTC

*: RDM: Default I/F B/W % : 75% [default] (max resv/bc0), 0% [default] (bc1)

Interface   MaxBW (bps)  MaxFlow (bps) Allocated (bps)      MaxSub (bps)
----------- ------------ ------------- -------------------- -------------
Gi0/0/0/0.519        750M          750M          35M (  4%)           0
Gi0/0/0/0.619        750M          750M            0 (  0%)           0
```

The above output shows that R3 can support up to 30 Mbps of reservations, but currently has no reservations actually in use. R4 can support up to 750 Mbps of reservations, and it currently has reservations of 35 Mbps and 50 Mbps on interfaces Gig1.24 and Gig1.45. XR1 can support up to 750 Mbps of reservations, and it currently has 35 Mbps reserved on Gig0/0/0/0.519. Note that in regular IOS when you issue the `ip rsvp bandwidth` command at the interface level, the default amount of reservable bandwidth is 75% of the configured `bandwidth` command. In IOS XR, there is no default and the amount of bandwidth available for reservations must be manually specified.

The details of all links through the entire TE topology, along with their current reservations and available bandwidth, can be verified at any point in the core network as follows.

CONTENTS

```
R5#show mpls traffic-eng topology 4.4.4.4

IGP Id: 0000.0000.0004.00, MPLS TE Id:4.4.4.4 Router Node  (isis  level-2) id 48
      link[0]: Broadcast, DR: 0000.0000.0002.02, nbr_node_id:41, gen:101
      frag_id: 0,  Intf Address: 20.2.4.4
      TE metric: 10, IGP metric: 10, attribute flags: 0x0
      SRLGs: None
       physical_bw: 1000000 (kbps), max_reservable_bw_global: 750000 (kbps)
      max_reservable_bw_sub: 0 (kbps)

                             Global Pool       Sub Pool
           Total Allocated   Reservable        Reservable
           BW (kbps)         BW (kbps)         BW (kbps)
           ---------------   -----------       ----------
    bw[0]:            0          750000                0
    bw[1]:            0          750000                0
    bw[2]:            0          750000                0
    bw[3]:            0          750000                0
    bw[4]:            0          750000                0
    bw[5]:            0          750000                0
    bw[6]:            0          750000                0
     bw[7]:        35000           715000                0

      link[1]: Broadcast, DR: 0000.0000.0003.02, nbr_node_id:53, gen:101
      frag_id: 0, Intf Address: 20.3.4.4
      TE metric: 10, IGP metric: 10, attribute flags: 0x0
      SRLGs: None
      physical_bw: 1000000 (kbps), max_reservable_bw_global: 750000 (kbps)
      max_reservable_bw_sub: 0 (kbps)

                             Global Pool       Sub Pool
           Total Allocated   Reservable        Reservable
           BW (kbps)         BW (kbps)         BW (kbps)
           ---------------   -----------       ----------
    bw[0]:            0          750000                0
    bw[1]:            0          750000                0
    bw[2]:            0          750000                0
    bw[3]:            0          750000                0
    bw[4]:            0          750000                0
    bw[5]:            0          750000                0
    bw[6]:            0          750000                0
    bw[7]:            0          750000                0

      link[2]: Broadcast, DR: 0000.0000.0004.03, nbr_node_id:46, gen:101
      frag_id: 0,  Intf Address: 20.4.5.4
      TE metric: 10, IGP metric: 10, attribute flags: 0x0
      SRLGs: None
      physical_bw: 1000000 (kbps), max_reservable_bw_global: 750000 (kbps)
      max_reservable_bw_sub: 0 (kbps)

                             Global Pool       Sub Pool
           Total Allocated   Reservable        Reservable
```

CONTENTS

```
          BW (kbps)         BW (kbps)         BW (kbps)
          ---------------   -----------       ----------
   bw[0]:            0          750000                0
   bw[1]:            0          750000                0
   bw[2]:            0          750000                0
   bw[3]:            0          750000                0
   bw[4]:            0          750000                0
   bw[5]:            0          750000                0
   bw[6]:            0          750000                0
    bw[7]:        50000           700000                0

      link[3]: Broadcast, DR: 0000.0000.0006.02, nbr_node_id:55, gen:101
      frag_id: 0, Intf Address: 20.4.6.4
      TE metric: 10, IGP metric: 10, attribute flags: 0x0
      SRLGs: None
      physical_bw: 1000000 (kbps), max_reservable_bw_global: 750000 (kbps)
      max_reservable_bw_sub: 0 (kbps)

                            Global Pool       Sub Pool
          Total Allocated   Reservable        Reservable
          BW (kbps)         BW (kbps)         BW (kbps)
          ---------------   -----------       ----------
   bw[0]:            0          750000                0
   bw[1]:            0          750000                0
   bw[2]:            0          750000                0
   bw[3]:            0          750000                0
   bw[4]:            0          750000                0
   bw[5]:            0          750000                0
   bw[6]:            0          750000                0
   bw[7]:            0          750000                0


R2#show mpls traffic-eng tunnels

P2P TUNNELS/LSPs:

Name: R2_t0                                (Tunnel0) Destination: 19.19.19.19
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 10, type dynamic (Basis for Setup, path weight 21)

  Config Parameters:
    Bandwidth: 50000     kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
   Metric Type: TE (default)
   AutoRoute: enabled  LockDown: disabled Loadshare: 50000 [40000] bw-based
   auto-bw: disabled
  Active Path Option Parameters:
    State: dynamic path option 10 is active
   BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled


  InLabel  :  -
  OutLabel : GigabitEthernet1.24, 18
  Next Hop : 20.2.4.4
```

```
  RSVP Signalling Info:
       Src 2.2.2.2, Dst 19.19.19.19, Tun_Id 0, Tun_Instance 7
    RSVP Path Info:
      My Address: 20.2.4.2
      Explicit Route: 20.2.4.4 20.4.5.4 20.4.5.5 20.5.19.5
                      20.5.19.19 19.19.19.19
      Record   Route:   NONE
      Tspec: ave rate=50000 kbits, burst=1000 bytes, peak rate=50000 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=50000 kbits, burst=1000 bytes, peak rate=50000 kbits
  History:
    Tunnel:
      Time since created: 36 minutes, 27 seconds
      Time since path change: 1 minutes, 11 seconds
      Number of LSP IDs (Tun_Instances) used: 7
    Current LSP: [ID: 7]
      Uptime: 1 minutes, 11 seconds
    Prior LSP: [ID: 4]
      ID: path option unknown
      Removal Trigger: signalling shutdown

 LSP Tunnel XR1_t0 is signalled, connection is up
  InLabel  : GigabitEthernet1.24, implicit-null
  Prev Hop : 20.2.4.4
  OutLabel :  -
  RSVP Signalling Info:
       Src 19.19.19.19, Dst 2.2.2.2, Tun_Id 0, Tun_Instance 4
    RSVP Path Info:
      My Address: 2.2.2.2
      Explicit Route:  NONE
      Record   Route:   NONE
      Tspec: ave rate=35000 kbits, burst=1000 bytes, peak rate=35000 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=35000 kbits, burst=1000 bytes, peak rate=35000 kbits

P2MP TUNNELS:

P2MP SUB-LSPS:
```

CONTENTS

We can also look at the link-manager, a data-structure maintained by the routers which keeps track of RSVP and MPLS-TE information. The link-manager performs admission control each time an RSVP message arrives at the router requesting resources. Additionally, it keep track of what information has been flooded or is pending to be flooded.

```
R2#show mpls traffic-eng link-management admission-control
System Information::
  P2P LSP:
    LSPS: 2 (Total Admitted:  60 , Total Deleted:  58 )
  P2MP subLSP:
    subLSPs: 0 (Total Admitted:  0 , Total Deleted:  0 )

P2P LSP:
  LSP ID               UP IF     DOWN IF  PRIORITY STATE          BW (kbps)
 2.2.2.2_1->19.19.19.1  -         Gi1.24   7/7      Resv Admitted   50000      RG
 19.19.19.19_2->2.2.2.  Gi1.24    -        7/7      Resv Admitted   0          U

P2MP subLSP:
  Src->Dst             UP IF     DOWN IF  PRIORITY STATE          BW (kbps)
```

```
R4#show mpls traffic-eng link-management admission-control
System Information::
  P2P LSP:
    LSPS: 2 (Total Admitted:  54 , Total Deleted:  52 )
  P2MP subLSP:
    subLSPs: 0 (Total Admitted:  0 , Total Deleted:  0 )

P2P LSP:
  LSP ID               UP IF     DOWN IF  PRIORITY STATE          BW (kbps)
 2.2.2.2_1->19.19.19.1  Gi1.24    Gi1.45   7/7      Resv Admitted   50000      RG
 19.19.19.19_2->2.2.2.  Gi1.45    Gi1.24   7/7      Resv Admitted   35000      RG

P2MP subLSP:
  Src->Dst             UP IF     DOWN IF  PRIORITY STATE          BW (kbps)
```

Note that the link-manager only keeps track of local interface data, unlike the TED or IGP database which is distributed.

```
R4#show mpls traffic-eng link-management interfaces
System Information::
    Links Count:         4
 Link ID::  Gi1.24 (20.2.4.4)
    Local Intfc ID:        6
    Link Status:
      SRLGs:                None
      Intfc Switching Capability Descriptors:
         Default:           Intfc Switching Cap psc1, Encoding ethernet
      Link Label Type:      Packet
       Physical Bandwidth:   1000000 kbits/sec
       Max Res Global BW:    750000 kbits/sec (reserved: 0% in, 5% out)
      Max Res Sub BW:       0 kbits/sec (reserved: 100% in, 100% out)
      MPLS TE Link State:   MPLS TE on, RSVP on, admin-up, flooded, allocated
      Inbound Admission:    reject-huge
      Outbound Admission:   allow-if-room
      Link MTU:             IP 1500, MPLS 1500
       Admin. Weight:        10 (IGP)
      IGP Neighbor Count:   1
       Neighbor:             ID 0000.0000.0002.02, IP 0.0.0.0 (Up)
    Flooding Status for each configured area [1]:
      IGP Area[1]:  isis  level-2:  flooded
 Link ID::  Gi1.34 (20.3.4.4)
    Local Intfc ID:        7
    Link Status:
      SRLGs:                None
      Intfc Switching Capability Descriptors:
         Default:           Intfc Switching Cap psc1, Encoding ethernet
      Link Label Type:      Packet
      Physical Bandwidth:   1000000 kbits/sec
      Max Res Global BW:    750000 kbits/sec (reserved: 0% in, 0% out)
      Max Res Sub BW:       0 kbits/sec (reserved: 100% in, 100% out)
      MPLS TE Link State:   MPLS TE on, RSVP on, admin-up, flooded
      Inbound Admission:    reject-huge
      Outbound Admission:   allow-if-room
      Link MTU:             IP 1500, MPLS 1500
      Admin. Weight:        10 (IGP)
      IGP Neighbor Count:   1
      Neighbor:             ID 0000.0000.0003.02, IP 0.0.0.0 (Up)
    Flooding Status for each configured area [1]:
      IGP Area[1]:  isis  level-2:  flooded
 Link ID::  Gi1.45 (20.4.5.4)
    Local Intfc ID:        8
    Link Status:
      SRLGs:                None
      Intfc Switching Capability Descriptors:
         Default:           Intfc Switching Cap psc1, Encoding ethernet
      Link Label Type:      Packet
       Physical Bandwidth:   1000000 kbits/sec
       Max Res Global BW:    750000 kbits/sec (reserved: 0% in, 7% out)
      Max Res Sub BW:       0 kbits/sec (reserved: 100% in, 100% out)
      MPLS TE Link State:   MPLS TE on, RSVP on, admin-up, flooded, allocated
      Inbound Admission:    reject-huge
```

```
      Outbound Admission:   allow-if-room
      Link MTU:             IP 1500, MPLS 1500
      Admin. Weight:        10 (IGP)
      IGP Neighbor Count:   1
      Neighbor:             ID 0000.0000.0004.03, IP 0.0.0.0 (Up)
    Flooding Status for each configured area [1]:
      IGP Area[1]:  isis  level-2:  flooded
Link ID::  Gi1.46 (20.4.6.4)
    Local Intfc ID:        9
    Link Status:
      SRLGs:                None
      Intfc Switching Capability Descriptors:
         Default:           Intfc Switching Cap psc1, Encoding ethernet
      Link Label Type:      Packet
      Physical Bandwidth:   1000000 kbits/sec
      Max Res Global BW:    750000 kbits/sec (reserved: 0% in, 0% out)
      Max Res Sub BW:       0 kbits/sec (reserved: 100% in, 100% out)
      MPLS TE Link State:   MPLS TE on, RSVP on, admin-up, flooded
      Inbound Admission:    reject-huge
      Outbound Admission:   allow-if-room
      Link MTU:             IP 1500, MPLS 1500
      Admin. Weight:        10 (IGP)
      IGP Neighbor Count:   1
      Neighbor:             ID 0000.0000.0006.02, IP 0.0.0.0 (Up)
    Flooding Status for each configured area [1]:
      IGP Area[1]:  isis  level-2:  flooded
```

CONTENTS

Like the OSPF or IS-IS databases, the TED output can show all routers in the topology at the same time, or it can show just an individual node by referencing its MPLS TE Router-ID, as seen above. In this case, we see that the node with Router-ID 4.4.4.4 (R4) has a link with address 20.4.5.4 (the link to R5), which has a total bandwidth of 1000 Mbps, a maximum reservable bandwidth of 750 Mbps (75% by default of the total bandwidth), 50 Mbps actually reserved, and 700 Mbps left over. The number 7 of the reservation refers to the TE tunnel's setup and hold priority that actually made the reservation. If a new reservation were to come in for 65 Mbps, with a higher setup priority (lower numerical value), the new tunnel could preempt the old tunnel's reservation. In this manner, the core of the network can offer prioritization of tunnel preference.

The final result of this configuration is that the tunnel from R2 to XR1 and the tunnel from XR1 to R2 avoid the links of R3 and R6, because they do not have enough available bandwidth for reservation. This can be seen in the traceroute outputs below.

```
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 20.2.4.4 [MPLS: Labels 18/16015 Exp 0] 9 msec 8 msec 6 msec
  3 20.4.5.5 [MPLS: Labels 17/16015 Exp 0] 25 msec 31 msec 31 msec
  4 20.5.19.19 21 msec 14 msec 15 msec
  5 10.19.20.20 16 msec *  10 msec



RP/0/0/CPU0:XR2#traceroute 1.1.1.1
Fri May 29 00:24:08.524 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 18/19 Exp 0] 29 msec  9 msec  9 msec
 3  20.4.5.4 [MPLS: Labels 16/19 Exp 0] 9 msec  0 msec  0 msec
 4  10.1.2.2 [MPLS: Label 19 Exp 0] 49 msec  19 msec  9 msec
 5  10.1.2.1 0 msec  *  0 msec
```

The disadvantage of the design of this specific scenario is that if any of the links of either R4 or R5 go down, both tunnels will go down, because there is not enough bandwidth for them to recalculate to use the paths via R3 or R6. The link failure and resulting tunnel teardown are shown below.

```
R2#debug mpls traffic-eng tunnel error
TE-SIG-HE: Tunnel0 [7]: path verification failed (unprotected) [Can't use link 20.5.19.5 on node 5.5.5.5]
R2#
TE-SIG-LM: tunnel path/reservation teardown failed : Tunnel not found (state may have been deleted already)
R2#
%LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed state to down


R5#config t
Enter configuration commands, one per line. End with CNTL/Z.
R5(config)#int Gig1
R5(config-if)#shut
R5(config-if)#
%LINK-5-CHANGED: Interface GigabitEthernet1, changed state to administratively down
%LINEPROTO-5-UPDOWN: Line protocol on Interface GigabitEthernet1, changed state to
down
R5(config-if)#
```

Because R2 has no alternate path in the topology with an available bandwidth of 50 Mbps, the tunnel cannot reroute around the failure.

If the bandwidth requirements were lowered, the tunnel would automatically set up again via the R3 and R6 path. We can validate this by looking at the following command, which displays what CSPF would compute given a particular input. There is no valid destination when 50 Mbps are requested - this is already known as the tunnel went down.

CONTENTS

```
R2#show mpls traffic-eng topology path destination 19.19.19.19 bandwidth 50000

Query Parameters:

  Destination: 19.19.19.19

     Bandwidth: 50000

   Priorities: 0 (setup), 0 (hold)

     Affinity: 0x0 (value), 0xFFFFFFFF (mask)

Query Results:

 % No matching path to destination, 19.19.19.19
```

If less bandwidth is requested, 10 Mbps for example, CSPF would be able to find a valid path:

```
R2#show mpls traffic-eng topology path destination 19.19.19.19 bandwidth 10000

Query Parameters:

  Destination: 19.19.19.19

     Bandwidth: 10000

   Priorities: 0 (setup), 0 (hold)

     Affinity: 0x0 (value), 0xFFFFFFFF (mask)

Query Results:

  Min Bandwidth Along Path: 0 (kbps)

  Max Bandwidth Along Path: 750000 (kbps)

  Hop  0: 20.2.4.2            : affinity 0x00000000, bandwidth 750000 (kbps)

  Hop  1: 20.2.4.4            : affinity 0x00000000, bandwidth 750000 (kbps)

  Hop  2: 20.4.6.4            : affinity 0x00000000, bandwidth 750000 (kbps)

  Hop  3: 20.4.6.6            : affinity 0x00000000, bandwidth 30000 (kbps)

  Hop  4: 20.6.19.6           : affinity 0x00000000, bandwidth 30000 (kbps)

  Hop  5: 20.6.19.19          : affinity 0x00000000, bandwidth 0 (kbps)

  Hop  6: 19.19.19.19
```

This can be further verified by actually requesting 10Mbps from the tunnel:

```
R2#debug ip rsvp signalling
RSVP signalling debugging is on

R2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#interface Tunnel0
R2(config-if)# tunnel mpls traffic-eng bandwidth 10000
R2(config-if)#end
R2#
 %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed state to up

R2#show mpls traffic-eng tunnels


P2P TUNNELS/LSPs:

Name: R2_t0                               (Tunnel0) Destination: 19.19.19.19
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 10, type dynamic (Basis for Setup, path weight 21)

  Config Parameters:
    Bandwidth: 10000    kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute: enabled  LockDown: disabled Loadshare: 10000 [200000] bw-based
    auto-bw: disabled
  Active Path Option Parameters:
    State: dynamic path option 10 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled



  InLabel  :  -
  OutLabel : GigabitEthernet1.24, 19
  Next Hop : 20.2.4.4
  RSVP Signalling Info:
       Src 2.2.2.2, Dst 19.19.19.19, Tun_Id 0, Tun_Instance 21
    RSVP Path Info:
      My Address: 20.2.4.2
      Explicit Route: 20.2.4.4 20.4.6.4 20.4.6.6 20.6.19.6
                      20.6.19.19 19.19.19.19
      Record   Route:   NONE
      Tspec: ave rate=10000 kbits, burst=1000 bytes, peak rate=10000 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=10000 kbits, burst=1000 bytes, peak rate=10000 kbits
  History:
    Tunnel:
      Time since created: 42 minutes, 45 seconds
      Time since path change: 42 seconds
      Number of LSP IDs (Tun_Instances) used: 21
    Current LSP: [ID: 21]
      Uptime: 42 seconds
    Prior LSP: [ID: 7]
```

```
      ID: path option unknown

      Removal Trigger: path verification failed


P2MP TUNNELS:


P2MP SUB-LSPS:
```

« MPLS TE Explicit Paths (/workbook/view/service-provider-v4/task/mpls-te-explicit-paths-Mjg4MQ%3D%3D) |
MPLS TE Metric Manipulation (/workbook/view/service-provider-v4/task/mpls-te-metric-manipulation-Mjg4Mw%3D%3D) »

CONTENTS