MPLS TE with Static Routing

« MPLS TE Metric Manipulation (/workbook/view/service-provider-v4/task/mpls-te-metric-manipulation-Mjg4Mw%3D%3D) | MPLS TE with Forwarding Adjacency (/workbook/view/service-provider-v4/task/mpls-te-with-forwarding-adjacency-MjkzMg%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 dynamic.
  - Configure a static route for XR1’s Loopback0 interface via the tunnel.
- 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 dynamic.
  - Configure a static route for R2’s Loopback0 interface via the tunnel.
- 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***

```
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 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
!
ip route 19.19.19.19 255.255.255.255 Tunnel0

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
 ip rsvp bandwidth
!
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
!
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

XR1:
interface tunnel-te0
 ipv4 unnumbered Loopback0
 destination 2.2.2.2
 path-option 1 dynamic
!
router static
 address-family ipv4 unicast
  2.2.2.2/32 tunnel-te0
!
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

In the previous examples, the MPLS TE tunnels of R2 and XR1 were advertised into the IGP process with the `autoroute announce` option. Because a TE tunnel is a unidirectional tunnel, there is no actual IGP adjacency that is established over it, but it can be used as a one-way link. This means that a TE tunnel can never receive packets, only send packets down the LSP towards the tail. For this reason, an IGP adjacency could never be established over a TE tunnel. `autoroute announce` essentially makes the tunnel interface seem as if it were directly connected to the tail end of the tunnel. Following this idea, any packets sent down the tunnel should be able to reach the tunnel tail, and any other devices behind the tail. To account for `autoroute announce`, the IGP installs a route in the RIB for the tunnel tail, and any nodes that are reachable behind the tail, with an outgoing interface of the TE tunnel instead of the underlying IGP network.

An alternative to this, as seen in this example, is to simply install a static route to the desired destination out the tunnel interface. All other verifications for this task are identical to the previous tasks, with the final result being that the customer traffic is routed over the TE tunnel in the core.

```
R2#show ip route 19.19.19.19
Routing entry for 19.19.19.19/32
  Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks:
  * directly connected, via Tunnel0
      Route metric is 0, traffic share count is 1

RP/0/0/CPU0:XR1#show route 2.2.2.2
Sat May 30 00:15:35.691 UTC

Routing entry for 2.2.2.2/32
  Known via "static", distance 1, metric 0 (connected)
  Installed May 30 00:07:23.514 for 00:08:12
  Routing Descriptor Blocks
    directly connected, via tunnel-te0
      Route metric is 0
  No advertising protos.

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 2 msec 1 msec
  2 20.2.4.4 [MPLS: Labels 16/16014 Exp 0] 9 msec 6 msec 5 msec
  3 20.4.5.5 [MPLS: Labels 16/16014 Exp 0] 26 msec 31 msec 31 msec
  4 20.5.19.19 21 msec 15 msec 14 msec
  5 10.19.20.20 16 msec *  11 msec

RP/0/0/CPU0:XR2#traceroute 1.1.1.1
Sat May 30 00:37:30.741 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 17/17 Exp 0] 29 msec  9 msec  0 msec
 3  20.4.5.4 [MPLS: Labels 17/17 Exp 0] 0 msec  0 msec  0 msec
 4  10.1.2.2 [MPLS: Label 17 Exp 0] 189 msec  9 msec  9 msec
 5  10.1.2.1 0 msec  *  0 msec
```

Note that Policy-Based Routing can also be use to send traffic through a TE tunnel.

« MPLS TE Metric Manipulation (/workbook/view/service-provider-v4/task/mpls-te-metric-manipulation-Mjg4Mw%3D%3D) | MPLS TE with Forwarding Adjacency (/workbook/view/service-provider-v4/task/mpls-te-with-forwarding-adjacency-MjkzMg%3D%3D) »

CONTENTS