RSVP TE Lab

Resource Reservation Protocol Traffic Engineering

Topology:

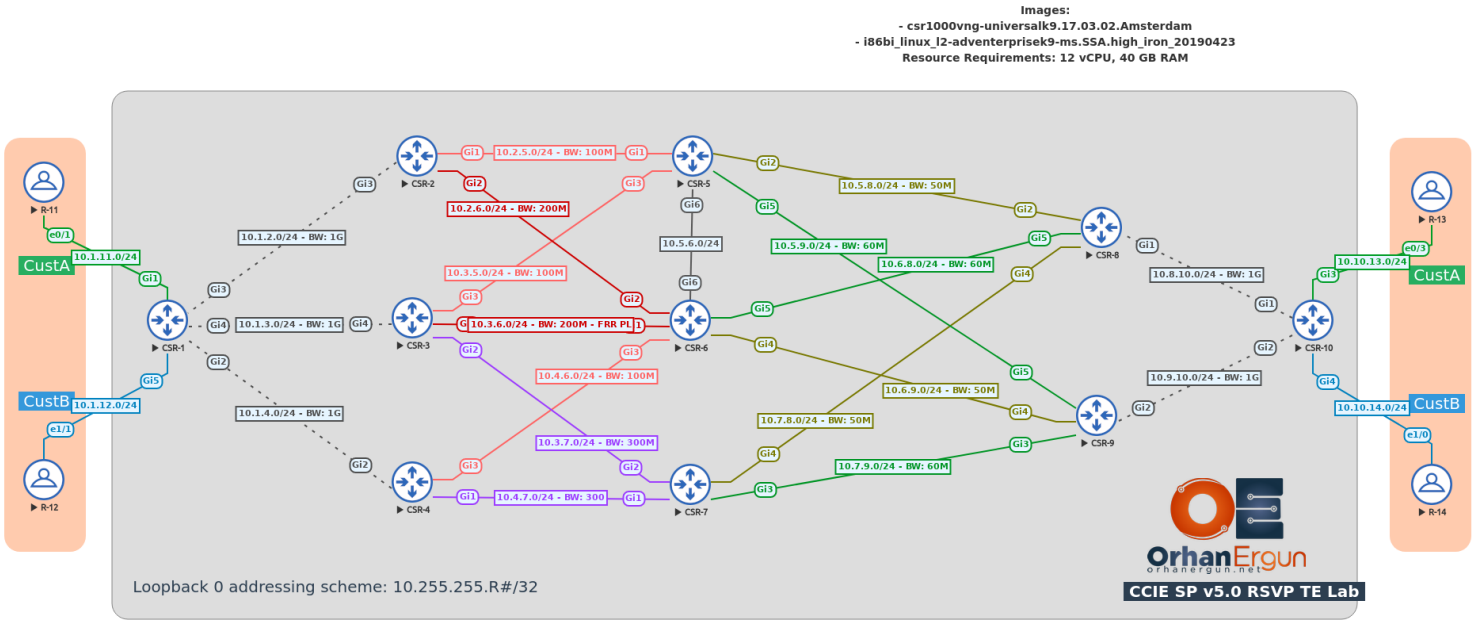

Task 01:

- Enable MPLS Traffic Engineering on all of the routers
- Set the IP RSVP bandwidth values on all Service Provider nodes (RSVP bandwidth sould use 100% of the interface bandwidth)

Solution:

What is Traffic Engineering and why we need it?

We have been using IGPs for many years, each IGP has it's own metric calculation formula and the metric considerations. For example EIGRP, by default considers Bandwidth and Delay of the interface and puts them in it's formula so as a result it gives us a composite metric.

OSPF uses the interface bandwidth and has Auto-Cost feature that the administrator can change it to any desired value, then the link metrics will be calculated based on that value (OSPF RFC does not contain: you need to use bandwidth as a metric, but all Router vendors are using bandwidth as the metric for OSPF).

When you run these IGPs, they are going to find all the best paths towards some specific destination and you don't have any control on the packet path (by default). If you want to alter these paths selections, you can make some changes to the metrics of some links and you will achieve some kind of Traffic Engineering. Or you can use Policy Based Routing (PBR) to selectively direct some specific traffic on some links.

But imagin if we had a way of doing all these stuff in a better way or even a better control and flexibility. Fortunatly there is always a way! RSVP TE can provide us a bunch of tools to do the traffic engineering more professional and easier compared to the PBR or Bandwidth Delay manipulation.

So far, We know that there are some protocols that can be used to generate and distribute MPLS Labels, such as TDP, LDP and SR. RSVP can also distribute MPLS Labels. There is no always need to have LDP when you are implementing MPLS, RSVP itself can also be used for the label generation and advertisement purposes.

In this lab we are using IOS-XE devices (CSR1000v) in the SP core, this device supports full featured RSVP, we could also use IOL or vIOS virtual platforms but they do not provide all the features that we need for complete all these tasks.

By default, RSVP TE is not enabled on IOS, IOS-XE devices, the first step is to enable RSVP and MPLS Traffic engineering tunnels, because: RSVP is going to create a virtual tunnel between the desired devices (Head-End: The Tunnel starting endpoint and Tail-End: The tunnel ending endpoint).

Let's enable MPLS Traffic engineering on all SP routers:

```
On all SP routers:

mpls traffic-eng tunnel
```

Under all the SP routers interfaces, we should also enab;e mpls traffic-eng tunnels (Please load the Initial configuration, they are already configured), as an example on CSR-1:

```
CSR-1:

interface range g2-4
 mpls traffic-eng tunnels
!
```

Whenever you enable MPLS Traffic Engineering tunnel on an interface:

By default 75% of the interface bandwidth could be reserved by the RSVP, as an example if you have a GigabitEthernet interface, 750 Mbit/s could be reserved by RSVP (By default), in order to allow RSVP to reserve 100% of the interface total bandwidth:

```
CSR-1:

interface range g2-4
 ip rsvp bandwidth percent 100
!
```

That is all you need to do (on all SP routers) in order to enable RSVP and MPLS TE tunnels.

But how can the routers get information about all these interface Traffic Engineering parameters or constraints? (for example reservable bandwidth of the links), OSPF and IS-IS can actually carry these kind of information (Link State Routing Protocols). Let's do the next task for this purpose.

Task 02:

- Enable MPLS Traffic Engineering on ISIS

Solution:

On all the SP routers:

```
CSR-1 - CSR-10:
router isis
 metric-style wide
 mpls traffic-eng router-id Loopback0
 mpls traffic-eng level-2
!
```

In IS-IS the metric-style should be set to wide in order to enable the requierd TLVs to support Traffic Engineering parameters or constraints.

Task 03:

- Create TE Tunnel 0 (on CSR-1)
- Tunnel destination must be 10.255.255.10
- This tunnel should use 1Mbps BW
- The path selection must be Dynamic
- The traffic destined to 10.254.0.10 should use this tunnel (Use Static Route)

Solution:

The tunnels are being created on the Head-End device (The starting point of the tunnel), in this example, CSR-1 is our Head-End device.

NOTE: Always remember that all these tunnels are Unidirectional, in the lab environment for simplicity and educational purpose we create the tunnels in a Unidirectional way, in real SP you need to also considedr the return traffic, so the same tunnel typically are being created at the other end of the tunnel (On CSR-10).

We can understand that: A router being Head-End or Tail-End depends on our point of view, if the traffic is going to enter the tunnel on CSR-1, that router is Head-End; if the traffic is going to enter the tunnel on CSR-10, that router is our Head-End and so on… .

Every tunnel must have a bandwidth requirement, it can be static BW or Dynamically calculated by router itself based on the traffic passing that tunnel at a specific amount of time. In this example we are using the first option: Statically setting the BW requirement of the tunnel.

When the tunnel is formed, that specified amount of bandwidth is going to be reserved on each of the routers along the path.

CSR-1:

```
interface Tunnel0
 description TE-10.254.0.10-Dynamic
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng path-selection metric igp
```

- In order to make the tunnel up and running, it must have an IP address, it can be any IP address, but typically it is defined as an unnummbered IP address of router's Looback interface.
- Tunnel mode must be changed to the mpls traffic-eng.
- Every tunnel should have a destination (in this example CSR-10 Loopback 0 IP address)
- By default Priority 7 7 is automatically added by IOS (We will discuss about setup and hold priorities on upcoming tasks)
- We specified the Bandwidth requirement for setting up this tunnel (in kbps)

There is also a path-option command, what does this command do? And what are the other options?

When we create an RSVP Traffic-Engineering tunnel on the Head-End router, the router is going to run an algorithm called CSPF, it is actually the SPF (Shortest Path First) algorithm but considers some parameters (Constraints) in shortest path calculation (For example Bandwidth requirements and available bandwidth of each link on the other routers in the topology.

In order to understand this let's take a look at an LSP information:

```
CSR-1(config)#do sh isis database CSR-6.00-00 verbose


Tag null:


IS-IS Level-2 LSP CSR-6.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd      ATT/P/OL
CSR-6.00-00           0x0000000A   0x5184                747/1198      0/0/0
  Auth:         Algorithm Text, Length: 9
  Area Address: 49.0000
  NLPID:        0xCC
  Router ID:    10.255.255.6
  Hostname: CSR-6
  Metric: 5          IS-Extended CSR-2.00
    Interface IP Address: 10.2.6.6
    Neighbor IP Address: 10.2.6.2
    Affinity: 0x00000000
    Admin. Weight: 3000
    Physical BW: 200000 kbits/sec
    Reservable Global Pool BW: 200000 kbits/sec
```

```
    Global Pool BW Unreserved:
      [0]:   200000 kbits/sec, [1]:   200000 kbits/sec
      [2]:   200000 kbits/sec, [3]:   200000 kbits/sec
      [4]:   200000 kbits/sec, [5]:   200000 kbits/sec
      [6]:   200000 kbits/sec, [7]:   200000 kbits/sec
    Affinity: 0x00000000
    Admin. Weight: 3000
    Physical LINK BW: 200000 kbits/sec
  Metric: 5          IS-Extended CSR-3.00
    Interface IP Address: 10.3.6.6
    Neighbor IP Address: 10.3.6.3
    Affinity: 0x00000000
    Admin. Weight: 3000
    Physical BW: 200000 kbits/sec
    Reservable Global Pool BW: 200000 kbits/sec
    Global Pool BW Unreserved:
      [0]:   200000 kbits/sec, [1]:   200000 kbits/sec
      [2]:   200000 kbits/sec, [3]:   200000 kbits/sec
      [4]:   200000 kbits/sec, [5]:   200000 kbits/sec
      [6]:   200000 kbits/sec, [7]:   200000 kbits/sec
    Affinity: 0x00000000
    Admin. Weight: 3000
    Physical LINK BW: 200000 kbits/sec
  Metric: 10         IS-Extended CSR-4.00
    Interface IP Address: 10.4.6.6
    Neighbor IP Address: 10.4.6.4
    Affinity: 0x00000000
    Admin. Weight: 2000
    Physical BW: 100000 kbits/sec
    Reservable Global Pool BW: 100000 kbits/sec
    Global Pool BW Unreserved:
      [0]:   100000 kbits/sec, [1]:   100000 kbits/sec
      [2]:   100000 kbits/sec, [3]:   100000 kbits/sec
      [4]:   100000 kbits/sec, [5]:   100000 kbits/sec
```

```
    [6]:   100000 kbits/sec, [7]:   100000 kbits/sec
  Affinity: 0x00000000
  Admin. Weight: 2000
  Physical LINK BW: 100000 kbits/sec
Metric: 16         IS-Extended CSR-8.00
  Interface IP Address: 10.6.8.6
  Neighbor IP Address: 10.6.8.8
  Affinity: 0x00000000
  Admin. Weight: 2600
  Physical BW: 60000 kbits/sec
  Reservable Global Pool BW: 60000 kbits/sec
  Global Pool BW Unreserved:
    [0]:    60000 kbits/sec, [1]:    60000 kbits/sec
    [2]:    60000 kbits/sec, [3]:    60000 kbits/sec
    [4]:    60000 kbits/sec, [5]:    60000 kbits/sec
    [6]:    60000 kbits/sec, [7]:    57000 kbits/sec
  Affinity: 0x00000000
  Admin. Weight: 2600
  Physical LINK BW: 60000 kbits/sec
Metric: 20         IS-Extended CSR-9.00
  Interface IP Address: 10.6.9.6
  Neighbor IP Address: 10.6.9.9
  Affinity: 0x00000000
  Admin. Weight: 2500
  Physical BW: 50000 kbits/sec
  Reservable Global Pool BW: 50000 kbits/sec
  Global Pool BW Unreserved:
    [0]:    50000 kbits/sec, [1]:    50000 kbits/sec
    [2]:    50000 kbits/sec, [3]:    50000 kbits/sec
    [4]:    50000 kbits/sec, [5]:    50000 kbits/sec
    [6]:    50000 kbits/sec, [7]:    49000 kbits/sec
  Affinity: 0x00000000
  Admin. Weight: 2500
  Physical LINK BW: 50000 kbits/sec
```

```
IP Address:   10.255.255.6
Metric: 0          IP 10.255.255.6/32
  Route Admin Tag: 1
  Prefix-attr: X:0 R:0 N:1
  Source Router ID: 10.255.255.6
Metric: 10         IP 10.4.6.0/24
  Prefix-attr: X:0 R:0 N:0
Metric: 5          IP 10.3.6.0/24
  Prefix-attr: X:0 R:0 N:0
Metric: 16         IP 10.6.8.0/24
  Prefix-attr: X:0 R:0 N:0
Metric: 20         IP 10.6.9.0/24
  Prefix-attr: X:0 R:0 N:0
Metric: 5          IP 10.2.6.0/24
  Prefix-attr: X:0 R:0 N:0
```

Interesting! Thanks to IS-IS, Each router has the detailed topological information all about the other routers neighbors, links, reservable bandwidth on that link, etc … .

CSPF (Constrained SPF) uses these information to find the Shortest path and do some actions to reserve the required bandwidth along the path.

You can statically configure the path as a static path option, or you can leave it to the router itself to find the best path based on IGP + Constraints or RSVP Weight + Constraints (We will discuss about Weights on upcoming tasks).

In this example we have used:

```
tunnel mpls traffic-eng path-option 1 dynamic
```

With putting Dynamic path option in this MPLS TE tunnel, we leave it to the router to decide about the path, It will take a look at all available bandwidth on all the links inside SP core to find a best one based on the IGP metric.

Wait a second! What is the difference between MPLS TE Dynamic path selection and legacy IGP best path selection?

Legacy IGP (For example IS-IS) best path selection was only based on the link metrics.

MPLS TE dynamic path selection by default is based on link metrics (if you don't set RSVP administrative weight) + constraints (such as available bandwidth etc …).

Verification:

```
CSR-1(config)#do sh ip int br | include Tunnel0
Tunnel0                10.255.255.1    YES TFTP    up                    up


CSR-1(config)#do sh mpls traffic-eng tunnels tunnel 0


Name: TE-10.254.0.10-Dynamic              (Tunnel0) Destination: 10.255.255.10
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 1, type dynamic (Basis for Setup, path weight 21)


  Config Parameters:
    Bandwidth: 1000     kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
    Metric Type: IGP (interface)
    Path-selection Tiebreaker:
      Global: not set   Tunnel Specific: not set   Effective: min-fill (default)
    Hop Limit: disabled
    Cost Limit: disabled
    Path-invalidation timeout: 10000 msec (default), Action: Tear
    AutoRoute: disabled LockDown: disabled Loadshare: 1000 [2000000] bw-based
    auto-bw: disabled
    Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
  Active Path Option Parameters:
    State: dynamic path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled
  Node Hop Count: 4


  InLabel  :  -
  OutLabel : GigabitEthernet2, 40027
  Next Hop : 10.1.4.4
  RSVP Signalling Info:
       Src 10.255.255.1, Dst 10.255.255.10, Tun_Id 0, Tun_Instance 48
    RSVP Path Info:
      My Address: 10.1.4.1
      Explicit Route: 10.1.4.4 10.4.7.7 10.7.9.9 10.9.10.10
```

```
                    10.255.255.10
    Record   Route:   NONE
    Tspec: ave rate=1000 kbits, burst=1000 bytes, peak rate=1000 kbits
  RSVP Resv Info:
    Record   Route:   NONE
    Fspec: ave rate=1000 kbits, burst=1000 bytes, peak rate=1000 kbits
Shortest Unconstrained Path Info:
  Path Weight: 21 (IGP)
  Explicit Route: 10.1.4.4 10.4.7.7 10.7.9.9 10.9.10.10
                  10.255.255.10


History:
  Tunnel:
    Time since created: 1 hours, 44 minutes
    Time since path change: 1 hours, 32 minutes
    Number of LSP IDs (Tun_Instances) used: 48
  Current LSP: [ID: 48]
    Uptime: 1 hours, 32 minutes
```

CSR-1 calculated the shortest path and decided to signal some routers along this path to reserve 1Mbps bandwidth for this tunnel:

`Explicit Route: 10.1.4.4 10.4.7.7 10.7.9.9 10.9.10.10 10.255.255.10`

This path includes: CSR-4 -> CSR-7 -> CSR-9 -> CSR-10.

The process is:

- CSR-1 sends a Path message to CSR-4
- CSR-4 sends a Path message to CSR-7
- CSR-7 sends a Path message to CSR-9
- CSR-9 sends a Path message to CSR-10

- CSR-10 sends a Reserve message to CSR-9 + a label for the tunnel (because of PHP, CSR-10 will advertise label 3 to CSR-9)

- CSR-9 sends a Reserve message to CSR-7 + a label

- CSR-7 sends a Reserve message to CSR-4 + a label
- CSR-4 sends a Reserve message to CSR-1 + a label

That simple! There is a hop-by-hop signaling in RSVP starting from the Head-End ending on the Tail-End (For the reply, the same thing happens in the reverse direction).

Now if you take a look at above MPLS TE Tunnel information, you can find the outgoing label of this tunnel traffic:

```
OutLabel : GigabitEthernet2, 40027
```

NOTE: The tunnel itself is just a tunnel! Nothing happens if you don't redirect any traffic inside this tunnel. There are multiple ways of directing the traffic into this tunnel. The simplest one is Static Routes:

```
CSR-1:
ip route 10.254.0.10 255.255.255.255 Tunnel0 name TE_Dynamic
!
```

Verification:

```
CSR-1(config)#do trace 10.254.0.10 source lo 0
Type escape sequence to abort.
Tracing the route to 10.254.0.10
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.4.4 [MPLS: Label 40027 Exp 0] 26 msec 2 msec 2 msec
  2 10.4.7.7 [MPLS: Label 70026 Exp 0] 3 msec 2 msec 2 msec
  3 10.7.9.9 [MPLS: Label 90017 Exp 0] 3 msec 2 msec 2 msec
  4 10.9.10.10 4 msec 2 msec 2 msec
```

Task 04:

- Create TE Tunnel 1
- Tunnel destination must be 10.255.255.10
- This tunnel should use 1Mbps
- The path selection must be Explicit (R3 -> R6 -> R8 -> R10)
- The traffic destined to 10.254.1.10 should use this tunnel (Use Static Route)

Solution:

In the previous task we have discussed about Dynamic Paths.

This time, we will take a look at the static options. You can define some Explicit path option and force the Head-End router to reserve an LSP on the exact path that administrator wants:

```
CSR-1:
ip explicit-path name R3-R6-R8-R10 enable
 index 10 next-address 10.1.3.3
 index 20 next-address 10.3.6.6
 index 30 next-address 10.6.8.8
 index 40 next-address 10.8.10.10
!
interface Tunnel1
 description TE-10.254.1.10-Explicit01
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name R3-R6-R8-R10
 tunnel mpls traffic-eng path-selection metric igp
!
ip route 10.254.1.10 255.255.255.255 Tunnel1 name TE_Explicit01
```

NOTE: The LSP for this TE Tunnel is ONLY going to be created IF these routers along the path have enough reservable bandwidth, otherwise the tunnel will not come up.

All the other commands inside Tunnel 1 are the same as previous task command, we simply skip explaining them.

Verification:

```
CSR-1(config)#$traffic-eng tunnels tunnel 1 | include Explicit Route|Label
  InLabel  :  -
  OutLabel : GigabitEthernet4, 30029
      Explicit Route: 10.1.3.3 10.3.6.6 10.6.8.8 10.8.10.10
    Explicit Route: 10.1.4.4 10.4.7.7 10.7.9.9 10.9.10.10


CSR-1(config)#do trace 10.254.1.10 source lo 0
Type escape sequence to abort.
Tracing the route to 10.254.1.10
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.3.3 [MPLS: Label 30029 Exp 0] 4 msec 2 msec 2 msec
  2 10.3.6.6 [MPLS: Label 60025 Exp 0] 3 msec 2 msec 2 msec
  3 10.6.8.8 [MPLS: Label 80029 Exp 0] 3 msec 3 msec 4 msec
  4 10.8.10.10 5 msec 3 msec 3 msec
```

NOTE: As I mentioned before, This is just a One-Way or Unidirectional Tunnel. You need to configure another tunnel on CSR-10 (Tail-End router) to make sure LSP is bidirectional.

Task 05:

- Create TE Tunnel 2
- Tunnel destination must be 10.255.255.10
- This tunnel should use 1Mbps
- The path selection must be Explicit (traffic should not traverse R6)
- The traffic destined to 10.254.2.10 should use this tunnel (Use Static Route)

Solution:

In this Task, we will discuss about the explicit path option with a single difference:

CSR-1 can consider any router along the path to signal the TE LSP but: CSR-6 should be excluded/avoided:

```
CSR-1:
ip explicit-path name NOT-R6 enable
 index 1 exclude-address 10.2.6.6
 index 2 exclude-address 10.3.6.6
 index 3 exclude-address 10.4.6.6
!
interface Tunnel2
 description TE-10.254.2.10-Explicit02
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name NOT-R6
 tunnel mpls traffic-eng path-selection metric igp
!
ip route 10.254.2.10 255.255.255.255 Tunnel2 name NOT-R6
!
```

All CSR-6 interface IP addresses are excluded from the path selection, Any other router could be considered as the possible middle router along the LSP.

Verification:

```
CSR-1(config)#do sh mpls traffic-eng tunnels tunnel 2 | include Route|Label
    AutoRoute: disabled LockDown: disabled Loadshare: 1000 [2000000] bw-based
  InLabel  :  -
  OutLabel : GigabitEthernet4, 30030
      Explicit Route: 10.1.3.3 10.3.7.7 10.7.9.9 10.9.10.10
      Record   Route:   NONE
      Record   Route:   NONE

CSR-1(config)#do trace 10.254.2.10 source lo 0
Type escape sequence to abort.
Tracing the route to 10.254.2.10
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.3.3 [MPLS: Label 30030 Exp 0] 12 msec 2 msec 2 msec
  2 10.3.7.7 [MPLS: Label 70025 Exp 0] 3 msec 2 msec 2 msec
  3 10.7.9.9 [MPLS: Label 90016 Exp 0] 2 msec 2 msec 2 msec
  4 10.9.10.10 3 msec 2 msec 2 msec
```

R6 links have been excluded in the path calculation.

Task 06:

- Create TE Tunnel 3
- Tunnel destination must be 10.255.255.10
- This tunnel should use 1Mbps
- The path selection must be Semi-Dynamic (Loose usage) (first, It should go to R3 and then it must traverse R8)
- The traffic destined to 10.254.3.10 should use this tunnel (Use Static Route)

Solution:

Again, this time we will also use the explicit path option with a single difference compared to the previous tasks:

- CSR-3 should be included in the path
- CSR-8 must be included in the path

We can realize that: Only these two routers are required to be considered in the Path creation, We DO NOT care what happens in the middle!

So what actually happens to the path in between these two devices? The answer is Simple:

- IGP is going to take care of it.

```
CSR-1:
ip explicit-path name LOOSE enable
 index 1 next-address 10.1.3.3
 index 2 next-address loose 10.255.255.8
!
interface Tunnel3
 description TE-10.254.3.10-Explicit03
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name LOOSE
 tunnel mpls traffic-eng path-selection metric igp
!
ip route 10.254.3.10 255.255.255.255 Tunnel3 name Explicit03
```

Verification:

```
CSR-1(config)#do sh mpls traffic-eng tunnels tunnel 3 | include Route|Label
    AutoRoute: disabled LockDown: disabled Loadshare: 1000 [2000000] bw-based
  InLabel  :  -
  OutLabel : GigabitEthernet4, 30050
     Explicit Route: 10.1.3.3 10.3.6.6 10.6.8.8 10.255.255.8
     Record   Route:   NONE
     Record   Route:   NONE
    Explicit Route: 10.1.4.4 10.4.7.7 10.7.9.9 10.9.10.10


CSR-1(config)#do trace 10.254.3.10 source lo 0
Type escape sequence to abort.
Tracing the route to 10.254.3.10
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.3.3 [MPLS: Label 30050 Exp 0] 12 msec 2 msec 2 msec
  2 10.3.6.6 [MPLS: Label 60046 Exp 0] 2 msec 2 msec 2 msec
  3 10.6.8.8 [MPLS: Label 80052 Exp 0] 3 msec 2 msec 2 msec
  4 10.8.10.10 3 msec 2 msec 2 msec
```

Task 07:

- Create TE Tunnel 4
- Tunnel destination must be 10.255.255.10
- The Bandwidth requirement for this tunnel must be Automatic
- The path selection must be Dynamic
- The Auto-Bandwidth measurement for all of the tunnels (even in the future) must be set to 10 minutes
- The traffic destined to 10.254.4.10 should use this tunnel (Use Static Route)

Solution:

So far, untill this task we have been using static bandwidth settings for the tunnel. Bandwidth can be also calculated automatically by the IOS itself. The router itself is going to track the amount of bandwidth (average bandwidth) that is being used (utilized) by a tunnel and dynamically change the bandwidth value based on this average value. It is common in the SP networks to use the dynamic bandwidth, because: no one can exactly predict what amount of bandwidth is going to pass a tunnel.

```
CSR-1:
mpls traffic-eng auto-bw timers frequency 600
!
interface Tunnel4
 description TE-10.254.4.10-Dynamic_AutoBW
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng path-selection metric igp
 tunnel mpls traffic-eng auto-bw min-bw 8000
ip route 10.254.3.10 255.255.255.255 Tunnel3 name Explicit03
```

Auto-BW measurement set to 10 minutes for all the tunnels that are in Auto-BW mode (existing and the future tunnels), minimum bandwidth could be also set (it is set to 8 mbps in this task).

Task 08:

- Create TE Tunnel 5
- Tunnel destination must be 10.255.255.10
- The Bandwidth requirement for this tunnel must be set to 1Mbps
- The path selection should be Dynamic
- For CSPF calculation for this tunnel, IGP metrics should be used
- The traffic destined to 10.254.5.10 should use this tunnel (Use Static Route)

Solution:

In the first tasks, I have mentioned about the CSPF calculation, There are two ways of CSPF calculation:

- Based on IGP Metrics + Constraints
- Based on RSVP Administrative Weight + Constraints

If you don't configure any RSVP Weights on the interfaces, The CSPF uses IGP metrics by default.

The network administrator for many reasons can set a weight on the links (for example The links with minimal delay, Jitter etc…).

```
CSR-1:
interface GigabitEthernet2
 mpls traffic-eng administrative-weight 1001
!
interface Tunnel5
 description TE-10.254.5.10-IGP_METRIC
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng path-selection metric igp
!
ip route 10.254.5.10 255.255.255.255 Tunnel5 name IGP_METRIC
```

The administrator is already set the RSVP Weight, but inside the tunnel, He/She specified that

CSPF should be forced to use IGP as the path selection metric type.

Verification:

```
CSR-1(config)#$traffic-eng tunnels tunnel 5 | include Route|Label|IGP
Name: TE-10.254.5.10-IGP_METRIC           (Tunnel5) Destination: 10.255.255.10
    Metric Type: IGP (interface)
    AutoRoute: disabled LockDown: disabled Loadshare: 1000 [2000000] bw-based
  InLabel  :  -
  OutLabel : GigabitEthernet2, 40029
      Explicit Route: 10.1.4.4 10.4.7.7 10.7.9.9 10.9.10.10
      Record   Route:   NONE
      Record   Route:   NONE
    Path Weight: 21 (IGP)
    Explicit Route: 10.1.4.4 10.4.7.7 10.7.9.9 10.9.10.10


CSR-1(config)#do trace 10.254.5.10
Type escape sequence to abort.
Tracing the route to 10.254.5.10
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.4.4 [MPLS: Label 40029 Exp 0] 7 msec 2 msec 1 msec
  2 10.4.7.7 [MPLS: Label 70028 Exp 0] 2 msec 2 msec 2 msec
  3 10.7.9.9 [MPLS: Label 90019 Exp 0] 3 msec 2 msec 2 msec
  4 10.9.10.10 3 msec 2 msec 2 msec
```

Task 09:

- Set appropriate MPLS TE Administrative Weight on all SP node's links
- Create TE Tunnel 6
- Tunnel destination must be 10.255.255.10
- The Bandwidth requirement for this tunnel must be set to 1Mbps
- The path selection should be Dynamic
- For CSPF calculation for this tunnel, TE metrics should be used
- The traffic destined to 10.254.6.10 should use this tunnel (Use Static Route)

Solution:

This task is exactly like the Task 08 with only as single difference: CSPF should be based on RSVP TE Administrative Weights:

```
CSR-1:
interface GigabitEthernet2
 mpls traffic-eng administrative-weight 1001
!
interface GigabitEthernet3
 mpls traffic-eng administrative-weight 1003
!
interface GigabitEthernet4
 mpls traffic-eng administrative-weight 1002
!
interface Tunnel6
 description TE-10.254.6.10-TE_AD_WEIGHT
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng path-selection metric te
!
ip route 10.254.6.10 255.255.255.255 Tunnel6 name TE_AD_WEIGHT
```

NOTE: MPLS TE Administrative Weight should be configured on all SP routers interfaces (Already set with the initial configs).

Verification:

```
CSR-1(config)#do sh mpls traffic-eng tunnels tunnel 6 | include TE

Name: TE-10.254.6.10-TE_AD_WEIGHT         (Tunnel6) Destination: 10.255.255.10

    Metric Type: TE (interface)

    Path Weight: 6503 (TE)


CSR-1(config)#do trace 10.254.6.10 source lo 0

Type escape sequence to abort.

Tracing the route to 10.254.6.10

VRF info: (vrf in name/id, vrf out name/id)

  1 10.1.4.4 [MPLS: Label 40026 Exp 0] 4 msec 2 msec 2 msec

  2 10.4.6.6 [MPLS: Label 60024 Exp 0] 2 msec 2 msec 2 msec

  3 10.6.9.9 [MPLS: Label 90015 Exp 0] 3 msec 1 msec 1 msec

  4 10.9.10.10 3 msec 2 msec 2 msec
```

For the CSPF calculations, IGP metrics will NOT be used anymore, only TE metrics + Constraints will be used instead.

Task 10:

- Create TE Tunnel 7
- Tunnel destination must be 10.255.255.10
- The Bandwidth requirement for this tunnel must be set to 1Mbps
- The path selection should be Dynamic
- For CSPF calculation for this tunnel, TE metrics should be used
- Traffic Coming from CustA-R-11 VLAN100 destined to 10.254.7.10 sould be forwarded over this tunnel

Solution:

With all the previous tasks, we have been using A static route to direct the traffic over the MPLS TE tunnel.

There are some other options like, let's configure one of them (PBR: Policy Based Routing):

```
CSR-1:
ip access-list extended MATCH-NET100
 10 permit ip 10.100.11.0 0.0.0.255 host 10.254.7.10
!
route-map TE_PBR permit 10
 match ip address MATCH-NET100
 set interface Tunnel7
!
interface GigabitEthernet1.100
 encapsulation dot1Q 100
 ip address 10.100.11.1 255.255.255.0
 ip policy route-map TE_PBR
!
```

CustA R-11 is sending the VLAN 100 traffic to CSR-11 (dot1q tagged traffic), The PE device (CSR-1) is pushing the traffic into Tunnel 7:

```
CSR-1(config)#do sh mpls traffic-eng tunnels tunnel 7 | include Route
    AutoRoute: disabled LockDown: disabled Loadshare: 1000 [2000000] bw-based
      Explicit Route: 10.1.3.3 10.3.5.5 10.5.8.8 10.8.10.10
CSR-1(config)#do sh ip route | include Tunnel7
CSR-1(config)#
```

```
CSR-1(config)#do sh mpls traffic-eng tunnels tunnel 7 | include Route
    AutoRoute: disabled LockDown: disabled Loadshare: 1000 [2000000] bw-based
      Explicit Route: 10.1.3.3 10.3.5.5 10.5.8.8 10.8.10.10
CSR-1(config)#do sh ip route | include Tunnel7
CSR-1(config)#


R-11(config)#ip route 10.254.7.10 255.255.255.255 10.100.11.1
R-11(config)#do trace 10.254.7.10
Type escape sequence to abort.
Tracing the route to 10.254.7.10
VRF info: (vrf in name/id, vrf out name/id)
  1 10.100.11.1 1 msec 0 msec 0 msec
  2 10.1.3.3 [MPLS: Label 30028 Exp 0] 2 msec 2 msec 2 msec
  3 10.3.5.5 [MPLS: Label 50027 Exp 0] 2 msec 2 msec 1 msec
  4 10.5.8.8 [MPLS: Label 80028 Exp 0] 3 msec 1 msec 2 msec
  5 10.8.10.10 3 msec 2 msec 2 msec
!


CSR-1(config)#do sh route-map
route-map TE_PBR, permit, sequence 10
  Match clauses:
    ip address (access-lists): MATCH-NET100
  Set clauses:
    interface Tunnel7
  Policy routing matches: 24 packets, 1104 bytes
```

Task 11:

- Create TE Tunnel 8
- Tunnel destination must be 10.255.255.6
- The Bandwidth requirement for this tunnel must be set to 1Mbps
- The path selection should be Static (R4 -> R7 -> R3 -> R6)
- An automatic route over this tunnel should be automatically anounced into the routing table
- Check UCMP towards 10.255.255.8 in the ISIS learned routes

Solution:

In this task we will discuss about another method of using MPLS TE Tunnel interfaces, You already know that you can direct the traffic into the tunnel using Static Routes and PBR.

This time we want to intruduce another method which is called Auto-Anounce. The tunnel and it's route can be advertised into the IGP and put into the Routing Table Automatically, before we configure the tunnel, Let me show you something:

```
CSR-1(config)#do sh ip route 10.5.8.0

Routing entry for 10.5.8.0/24

  Known via "isis", distance 115, metric 42, type level-2

  Redistributing via isis

  Last update from 10.1.2.2 on GigabitEthernet3, 00:00:16 ago

  Routing Descriptor Blocks:

    10.1.4.4, from 10.255.255.8, 00:00:16 ago, via GigabitEthernet2

      Route metric is 42, traffic share count is 1

  * 10.1.3.3, from 10.255.255.8, 00:00:16 ago, via GigabitEthernet4

      Route metric is 42, traffic share count is 1

    10.1.2.2, from 10.255.255.8, 00:00:16 ago, via GigabitEthernet3

      Route metric is 42, traffic share count is 1
```

CSR-1 can reach to the 10.5.8.0 (on CSR-8) by forwarding the packets to the CSR-2 and CSR-3 and CSR-4, How did CSR-1 learn about this path? Using IS-IS learned routes.

We can create an MPLS TE Tunnel and anounce it to the IGP as well, so the IS-IS can consider this tunnel as a new link and path and import it to the RIB:

```
CSR-1:
ip explicit-path name R4-R7-R3-R6 enable
 index 1 next-address 10.1.4.4
 index 2 next-address 10.4.7.7
 index 3 next-address 10.3.7.3
 index 4 next-address 10.3.6.6
!
interface Tunnel8
 description TE-AutoRoute
 ip unnumbered Loopback0
 no shutdown
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.6
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name R4-R7-R3-R6
!
```

Verification:

```
CSR-1(config)#do sh ip route 10.5.8.0
Routing entry for 10.5.8.0/24
  Known via "isis", distance 115, metric 42, type level-2
  Redistributing via isis
  Last update from 10.255.255.6 on Tunnel8, 00:00:06 ago
  Routing Descriptor Blocks:
    10.255.255.6, from 10.255.255.8, 00:00:06 ago, via Tunnel8
      Route metric is 42, traffic share count is 1
  * 10.1.4.4, from 10.255.255.8, 00:00:06 ago, via GigabitEthernet2
      Route metric is 42, traffic share count is 1
    10.1.3.3, from 10.255.255.8, 00:00:06 ago, via GigabitEthernet4
      Route metric is 42, traffic share count is 1
```

We can achieve some kind of UCMP (Unequal Cost Multi-pathing) with even IS-IS or OSPF behind the scenes thanks to MPLS TE tunnels and Auto-Announce feature (not exactly the same as EIGRP UCMP but anyways we are happy with it ☺).

The metric of the route is 42, it's based on the CSPF calculated metric considering IGP metric.

```
CSR-1(config)#do sh isis topology

Tag null:

IS-IS TID 0 paths to level-2 routers
System Id            Metric     Next-Hop            Interface   SNPA
CSR-1                --
CSR-2                1          CSR-2               Gi3         5000.0002.0002
CSR-3                1          CSR-3               Gi4         5000.0003.0003
CSR-4                1          CSR-4               Gi2         5000.0004.0001
CSR-5                11         CSR-2               Gi3         5000.0002.0002
                                CSR-3               Gi4         5000.0003.0003
CSR-6                6          CSR-6               Tu8         *MPLS TE-Tunnel
CSR-7                4          CSR-3               Gi4         5000.0003.0003
                                CSR-4               Gi2         5000.0004.0001
CSR-8                22         CSR-3               Gi4         5000.0003.0003
                                CSR-4               Gi2         5000.0004.0001
                                CSR-6               Tu8         *MPLS TE-Tunnel
CSR-9                20         CSR-3               Gi4         5000.0003.0003
                                CSR-4               Gi2         5000.0004.0001
CSR-10               21         CSR-3               Gi4         5000.0003.0003
                                CSR-4               Gi2         5000.0004.0001
```

Interesting! IS-IS topology show us that there is a direct link between CSR-1 and two routers (CSR-6 and CSR-8), but in the physical topology there is no physical link between these routers. The link is an MPLS TE Tunnel.

Task 12:

- Create TE Tunnel 38 on R3 and R8
- This tunnel should be MPLS TE forwarding-Adjacency (use Lo0 as the tunnel destination)
- The Bandwidth requirement for this tunnel must be set to 1Mbps
- The path selection should be Dynamic
- Traffic sourced from R1 Loopback0 destined to R10 loopback0 sould go to R3 and enter this tunnel

Solution:

Things are getting interesting task by task! ☺

Let's create the tunnels on R3 and R8, then I will explain what happens with Forwarding-Adjacency type of tunnel in the verification section:

```
CSR-3:
interface Tunnel38
 description TE-FW_Adj
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.8
 tunnel mpls traffic-eng forwarding-adjacency
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 dynamic
!
CSR-8:
interface Tunnel38
 description TE-FW_Adj
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.3
 tunnel mpls traffic-eng forwarding-adjacency
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 dynamic
!
```

This kind of Tunnel must be configured in a Bidirectional form otherwise it won't work.

What FA means?

This is something like Auto-Route Anounce, but their difference is:

- Forwarding Adjacency is like forming a direct link between two routers, it will be anounced into the IS-IS or OSPF, and it is two-way tunnel (Bidirectional)
- Auto-Route Anounce is putting a One-Way tunnel into the IS-IS or OSPF, and it can have Multiple Dynamic endpoints based on the Metric (If it's metric is better than the Physical Links metric in order to reach some prefixes, the router can use it as a exit interface or if it's metric is equal to other next-hops, the router can do Load Sharing).

Verification:

```
CSR-3#show isis topology



IS-IS TID 0 paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
CSR-1                1          CSR-1                Gi4         5000.0001.0003
CSR-2                2          CSR-1                Gi4         5000.0001.0003
CSR-3                --
CSR-4                2          CSR-1                Gi4         5000.0001.0003
CSR-5                10         CSR-5                Gi3         5000.0005.0002
CSR-6                5          CSR-6                Gi1         5000.0006.0000
CSR-7                3          CSR-7                Gi2         5000.0007.0001
CSR-8                10         CSR-8                Tu38        *MPLS TE-Tunnel
CSR-9                12         CSR-8                Tu38        *MPLS TE-Tunnel
CSR-10               11         CSR-8                Tu38        *MPLS TE-Tunnel

CSR-3#show isis data CSR-8.00-00  detail



IS-IS Level-2 LSP CSR-8.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd      ATT/P/OL
CSR-8.00-00           0x00000019   0x6290                1092/1198      0/0/0
  Auth:         Algorithm Text, Length: 9
  Area Address: 49.0000
```

```
  NLPID:        0xCC
  Router ID:    10.255.255.8
  Hostname: CSR-8
  Metric: 20         IS-Extended CSR-5.00
  Metric: 16         IS-Extended CSR-6.00
  Metric: 20         IS-Extended CSR-7.00
  Metric: 1          IS-Extended CSR-10.00
  Metric: 10         IS-Extended CSR-3.00
  IP Address:   10.255.255.8
  Metric: 0          IP 10.255.255.8/32
  Metric: 20         IP 10.5.8.0/24
  Metric: 16         IP 10.6.8.0/24
  Metric: 20         IP 10.7.8.0/24
  Metric: 1          IP 10.8.10.0/24

IS-IS Level-2 LSP CSR-3.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd      ATT/P/OL
CSR-3.00-00         * 0x0000001E   0x49A5                1009/*         0/0/0
  Auth:         Algorithm Text, Length: 9
  Area Address: 49.0000
  NLPID:        0xCC
  Router ID:    10.255.255.3
  Hostname: CSR-3
  Metric: 1          IS-Extended CSR-1.00
  Metric: 5          IS-Extended CSR-6.00
  Metric: 10         IS-Extended CSR-5.00
  Metric: 3          IS-Extended CSR-7.00
  Metric: 10         IS-Extended CSR-8.00
  IP Address:   10.255.255.3
  Metric: 0          IP 10.255.255.3/32
  Metric: 10         IP 10.3.5.0/24
  Metric: 3          IP 10.3.7.0/24
  Metric: 5          IP 10.3.6.0/24
  Metric: 1          IP 10.1.3.0/24
```

```
CSR-3#show isis neighbors

System Id       Type Interface     IP Address      State Holdtime Circuit Id
CSR-1           L2   Gi4           10.1.3.1        UP    25       04
CSR-5           L2   Gi3           10.3.5.5        UP    28       03
CSR-6           L2   Gi1           10.3.6.6        UP    22       01
CSR-7           L2   Gi2           10.3.7.7        UP    29       02

CSR-3#show ip route isis | begin Gatewa
Gateway of last resort is not set

      10.0.0.0/8 is variably subnetted, 52 subnets, 2 masks
i L2     10.1.2.0/24 [115/2] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.1.4.0/24 [115/2] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.1.11.0/24 [115/1] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.2.5.0/24 [115/12] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.2.6.0/24 [115/7] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.4.6.0/24 [115/12] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.4.7.0/24 [115/5] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.5.8.0/24 [115/30] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.5.9.0/24 [115/26] via 10.3.5.5, 04:22:05, GigabitEthernet3
i L2     10.6.8.0/24 [115/21] via 10.3.6.6, 04:22:05, GigabitEthernet1
i L2     10.6.9.0/24 [115/25] via 10.3.6.6, 04:22:05, GigabitEthernet1
i L2     10.7.8.0/24 [115/23] via 10.3.7.7, 04:22:05, GigabitEthernet2
i L2     10.7.9.0/24 [115/19] via 10.3.7.7, 04:22:05, GigabitEthernet2
i L2     10.8.10.0/24 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.9.10.0/24 [115/12] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.100.11.0/24 [115/1] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.101.101.1/32 [115/1] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.101.101.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.102.102.1/32 [115/1] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.102.102.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.0.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.1.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.2.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
```

```
i L2     10.254.3.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.4.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.5.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.6.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.7.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.8.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.9.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.10.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.11.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.12.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.254.254.4/32 [115/2] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.255.255.1/32 [115/1] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.255.255.2/32 [115/2] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.255.255.4/32 [115/2] via 10.1.3.1, 04:22:05, GigabitEthernet4
i L2     10.255.255.5/32 [115/10] via 10.3.5.5, 04:22:05, GigabitEthernet3
i L2     10.255.255.6/32 [115/5] via 10.3.6.6, 04:22:05, GigabitEthernet1
i L2     10.255.255.7/32 [115/3] via 10.3.7.7, 04:22:05, GigabitEthernet2
i L2     10.255.255.8/32 [115/10] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.255.255.9/32 [115/12] via 10.255.255.8, 00:19:20, Tunnel38
i L2     10.255.255.10/32 [115/11] via 10.255.255.8, 00:19:20, Tunnel38
```

There is no IS-IS neighborship between CSR-3 and CSR-8 but we can find a direct link with the metric of 10 between these two routers in the IS-IS database.

Task 13:

- Create Tunnel 101 and 102 on R1
- The tunnel destination should be 10.255.255.10
- Tunnel 101 sould use expilicit path of (R1 -> R2 -> R5 -> R8 -> R10)
- Tunnel 102 sould use expilicit path of (R1 -> R4 -> R7 -> R9 -> R10)
- The Bandwidth requirement for this tunnel must be set to 1Mbps
- CustA VLAN 101 traffic should be forwarded over Tunnel 101
- CustB VLAN 102 traffic should be forwarded over Tunnel 102

PBTS (Policy Based Tunnel Selection)

Solution:

You are now familiar with different methods of directing the traffic inside MPLS TE Tunnels:

Using Static Routes, PBR.

Let's learn another method.

This time, CSR-1 (PE device) is providing MPLS L3VPN service to the customers.

We want to put the Customer traffic into a specific TE Tunnels. It's common in Service Providers to provide some services to the customer based on the SLA and the amount of payment they are doing! For example provide MPLS L3VPn GOLD, Silver, Bronze etc… .

First of all, let's create the requierd Tunnels:

```
CSR-1:
ip explicit-path name R1-R2-R5-R8-R10 enable
 index 1 next-address loose 10.255.255.2
 index 2 next-address loose 10.255.255.5
 index 3 next-address loose 10.255.255.8
 index 4 next-address loose 10.255.255.10
ip explicit-path name R1-R4-R7-R9-R10 enable
 index 1 next-address loose 10.255.255.4
 index 2 next-address loose 10.255.255.7
 index 3 next-address loose 10.255.255.9
 index 4 next-address loose 10.255.255.10
!
```

```
interface Tunnel101
 description TE-L3VPN_PerVRF_PE_PE_CustA
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name R1-R2-R5-R8-R10
!
interface Tunnel102
 description TE-L3VPN_PerVRF_PE_PE_CustB
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name R1-R4-R7-R9-R10
!
ip route 10.101.101.10 255.255.255.255 Tunnel101 name TE-L3VPN_PerVRF_PE_PE_CustA
ip route 10.102.102.10 255.255.255.255 Tunnel102 name TE-L3VPN_PerVRF_PE_PE_CustB
```

The tunnel destination is 10.255.255.10 but the static routes are towards 10.101.101.10 and 10.102.102.10!

CSR-1 and CSR-10 (PE routers) are forming MP-BGP VPNv4 Unicast neighborship using their Loopback 0 interfaces, so the next-hop of the VPNv4 routes would be 10.255.255.1 and 10.255.255.10, But why we are Directing the Customer traffic with the next-hop of 10.101.101.10 and 10.102.102.10 to those two tunnels?

This is a technique we can use along with the L3VPNs.

Let's check the CSR-1 configuration:

```
CSR-1:
router bgp 10000
address-family vpnv4
  neighbor 10.255.255.10 activate
  neighbor 10.255.255.10 send-community extended
exit-address-family
```

We can selectively change the VPNv4 routes next-hops on the remote PE:

```
CSR-10:
interface Loopback101
 ip address 10.101.101.10 255.255.255.255
!
interface Loopback102
 ip address 10.102.102.10 255.255.255.255
!
router bgp 10000
address-family vpnv4
  neighbor 10.255.255.1 activate
  neighbor 10.255.255.1 send-community extended
 exit-address-family
!
vrf definition CustA
 !
 address-family ipv4
  bgp next-hop Loopback101
 exit-address-family
!
vrf definition CustB
 !
 address-family ipv4
  bgp next-hop Loopback102
 exit-address-family
!
```

That's all you need to do, create two loopbacks on the remote PE device and change the Customer routes next-hops to those loopbacks.

Verification:

```
CSR-10(config)#do clear bgp vpnv4 uni * so


CSR-1(config)#do sh bgp vpnv4 uni all summ | include 10.255.255.10
10.255.255.10    4        10000     348     344       22    0    0 04:59:15        6
```

```
CSR-1(config)#do sh bgp vpnv4 uni all | begin Net
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf CustA)
 *>   10.101.1.0/24    0.0.0.0                  100         32768 ?
 *>i  10.101.10.0/24   10.101.101.10            100    100      0 ?
 *>   11.101.101.0/24  10.101.1.11               0             0 65001 i
 *>i  13.101.101.0/24  10.101.101.10             0    100      0 65002 i
 *>i  169.254.254.13/32
                       10.101.101.10             0    100      0 65002 i
Route Distinguisher: 1:2 (default for vrf CustB)
 *>   10.102.1.0/24    0.0.0.0                  100         32768 ?
 *>i  10.102.10.0/24   10.102.102.10            100    100      0 ?
 *>i  14.102.102.14/32 10.102.102.10             0    100      0 65004 i
 *>i  169.254.254.14/32
                       10.102.102.10             0    100      0 65004 i


CSR-1(config)#do sh ip route 10.101.101.10
Routing entry for 10.101.101.10/32
  Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks:
  * directly connected, via Tunnel101
      Route metric is 0, traffic share count is 1
CSR-1(config)#do sh ip route 10.102.102.10
Routing entry for 10.102.102.10/32
  Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks:
  * directly connected, via Tunnel102
      Route metric is 0, traffic share count is 1
```

Task 14:

- Provide VPWS (with XConnect) service to CustA VLAN 1113
- Configure Tunnel 1113 on R1
- Tunnel destination must be the lo0 address of the Remote PE device
- Path selection should be Expilicit (R1 -> R2 -> R6 -> R8 -> R10)
- The VC traffic must prefer this tunnel as it's primary path

Solution:

In the Task 13 we discussed about how we can direct the L3VPN customer traffic into a specific TE Tunnel, this time we will do the same thing but for L2VPN: VPWS service.

The configuration is straight forward and easy:

```
CSR-1:


interface Tunnel1113
 description TE-VPWS
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name VPWS_R1-R2-R6-R8-R10
!
ip explicit-path name VPWS_R1-R2-R6-R8-R10 enable
 index 1 next-address 10.1.2.2
 index 2 next-address 10.2.6.6
 index 3 next-address 10.6.8.8
 index 4 next-address 10.8.10.10
!
pseudowire-class R11-R13-PW
 encapsulation mpls
 control-word
 preferred-path interface Tunnel1113
!
```

```
interface GigabitEthernet1.1113
 encapsulation dot1Q 1113
 xconnect 10.255.255.10 1113 encapsulation mpls pw-class R11-R13-PW
!

CSR-10:
interface GigabitEthernet3.1113
 encapsulation dot1Q 1113
 xconnect 10.255.255.1 1113 encapsulation mpls
!
```

The pseudowire will prefer TE Tunnel 1113 in order to encapsulate the customer traffic and sent it to the remote PE device.

Task 15:

- Configure Tunnel 15 on R1 and R5
- Tunnel destination must be the lo0 address of the Remote PE device
- Path selection should be Expilicit (R1 -> R2 -> R5)
- CustA L3VPN traffic sould enter this PE to P MPLS TE tunnel and must have it's own LDP label (use Static route)

Solution:

In task 14 we created a TE tunnel between two PE devices. We can also create a TE tunnel between PE and core (P) device. In this scenario we are trying to send the traffic using some expilicit routers in between.

What about P to remote PE traffic? We simply do not care about how P device is going to direct the traffic towards the remote PE. The only thing the we do care is PE to P traffic path:

```
CSR-1:
ip explicit-path name PE_P_L3VPN_R1-R2-R5 enable
 index 1 next-address loose 10.255.255.2
 index 2 next-address loose 10.255.255.5
!
interface Tunnel15
 description TE-PE_To_P_L3VPN
 ip unnumbered Loopback0
 mpls ip
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.5
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name PE_P_L3VPN_R1-R2-R5
!
ip route 10.101.101.10 255.255.255.255 Tunnel15 name PE_P_L3VPN_R1-R2-R5
!
```

This time, inside the tunnel interface we also enable MPLS IP. So CSR-1 and CSR-5 are going to become LDP neighbors (Using Targeted LDP) and they can exchange Labels for this tunnel.

```
CSR-5:
ip explicit-path name PE_P_L3VPN_R5-R2-R1 enable
 index 1 next-address loose 10.255.255.2
 index 2 next-address loose 10.255.255.1
!
interface Tunnel15
 description TE-PE_To_P_L3VPN
 ip unnumbered Loopback0
 mpls ip
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.1
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name PE_P_L3VPN_R5-R2-R1
!
```

Verification:

```
CSR-1(config)#int tunnel 101
CSR-1(config-if)#shu
CSR-1(config-if)#
*Apr 14 23:53:47.628: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel101, changed state to down
*Apr 14 23:53:47.630: %LINK-5-CHANGED: Interface Tunnel101, changed state to administratively down
CSR-1(config-if)#do sh ip route 10.101.101.10
Routing entry for 10.101.101.10/32
  Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks:
  * directly connected, via Tunnel15
      Route metric is 0, traffic share count is 1
CSR-1(config)#do sh ip cef 10.101.101.10
10.101.101.10/32
  attached to Tunnel15 label 50041-(local:10058)
CSR-1(config)#do sh ip cef vrf CustA 13.101.101.0
13.101.101.0/24
  nexthop 10.101.101.10 Tunnel15 label 50041-(local:10058) 1000032
```

The customer traffic enters this Tunnel and it will have two labels (Inner Label or the VPN Label learned from MP-BGP which is 1000032 and an LSP label which is 50041.

Task 16:

- Color the links between R1, R2, R5, R8, R10 with the value of 0x00000001
- Color the links between R1, R4, R7, R9, R10 with the value of 0x00000002
- Create Two tunnels with dynamic path selection on R1, named Tunnel 10 and Tunnel 11
- The tunnel destination on both of the tunnel interfaces must be 10.255.255.10
- Tunnel 10 should prefer links with 0x00000001 color
- Tunnel 11 should prefer links with 0x00000002 color

Solution:

With RSVP TE, we can also have another constraint called Affinity. Using Affinity we can put some color on the links on each router. For example we can set some color for the links with lower delay, and later on we force the Head-End router to request for an LSP that includes those links (Those colored links with Low latency).

```
CSR-1:
int g3
 mpls traffic-eng attribute-flags 0x00000001
!
CSR-2:
int range g1-3
 mpls traffic-eng attribute-flags 0x00000001
!
CSR-5:
int range g1-2
 mpls traffic-eng attribute-flags 0x00000001
!
CSR-8:
int range g1-2
 mpls traffic-eng attribute-flags 0x00000001
!
CSR-10:
int g1
 mpls traffic-eng attribute-flags 0x00000001
!
```

CSR-1:

```
int g2
 mpls traffic-eng attribute-flags 0x00000002
!
```

CSR-4:

```
int range g1-2
 mpls traffic-eng attribute-flags 0x00000002
!
```

CSR-7:

```
int range g1, g3
 mpls traffic-eng attribute-flags 0x00000002
!
```

CSR-9:

```
int range g2-3
 mpls traffic-eng attribute-flags 0x00000002
!
```

CSR-10:

```
int g2
 mpls traffic-eng attribute-flags 0x00000002
!
```

CSR-1:

```
interface Tunnel10
 description TE-10.254.10.10-Affinity_0x00000001
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng affinity 0x00000001
 tunnel mpls traffic-eng path-option 1 dynamic
!
interface Tunnel11
 description TE-10.254.10.10-Affinity_0x00000002
 ip unnumbered Loopback0
```

```
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.10
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng affinity 0x00000002
 tunnel mpls traffic-eng path-option 1 dynamic
```

Tunnel 10 will request for an LSP along the path that has affinity 0x00000001 and Tunnel 11 do the same thing for affinity 0x00000002:

Verification:

```
CSR-1(config)#do sh isis database CSR-2.00-00 verbo


Tag null:


IS-IS Level-2 LSP CSR-2.00-00
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime/Rcvd      ATT/P/OL
CSR-2.00-00           0x00000025   0x89A7                1001/1199      0/0/0
  Auth:         Algorithm Text, Length: 9
  Area Address: 49.0000
  NLPID:        0xCC
  Router ID:    10.255.255.2
  Hostname: CSR-2
  Metric: 1          IS-Extended CSR-1.00
    Interface IP Address: 10.1.2.2
    Neighbor IP Address: 10.1.2.1
    Affinity: 0x00000001
    Admin. Weight: 1003
    Physical BW: 1000000 kbits/sec
    Reservable Global Pool BW: 1000000 kbits/sec
    Global Pool BW Unreserved:
      [0]:  1000000 kbits/sec, [1]:  1000000 kbits/sec
      [2]:  1000000 kbits/sec, [3]:  1000000 kbits/sec
      [4]:  1000000 kbits/sec, [5]:  1000000 kbits/sec
      [6]:  1000000 kbits/sec, [7]:   999000 kbits/sec
    Affinity: 0x00000001
```

```
    Admin. Weight: 1003
    Physical LINK BW: 1000000 kbits/sec
  Metric: 5          IS-Extended CSR-6.00
    Interface IP Address: 10.2.6.2
    Neighbor IP Address: 10.2.6.6
    Affinity: 0x00000000
    Admin. Weight: 3000
    Physical BW: 200000 kbits/sec
    Reservable Global Pool BW: 200000 kbits/sec
    Global Pool BW Unreserved:
      [0]:   200000 kbits/sec, [1]:   200000 kbits/sec
      [2]:   200000 kbits/sec, [3]:   200000 kbits/sec
      [4]:   200000 kbits/sec, [5]:   200000 kbits/sec
      [6]:   200000 kbits/sec, [7]:   200000 kbits/sec
    Affinity: 0x00000000
    Admin. Weight: 3000
    Physical LINK BW: 200000 kbits/sec
  Metric: 10         IS-Extended CSR-5.00
    Interface IP Address: 10.2.5.2
    Neighbor IP Address: 10.2.5.5
    Affinity: 0x00000001
    Admin. Weight: 2000
    Physical BW: 100000 kbits/sec
    Reservable Global Pool BW: 100000 kbits/sec
    Global Pool BW Unreserved:
      [0]:   100000 kbits/sec, [1]:   100000 kbits/sec
      [2]:   100000 kbits/sec, [3]:   100000 kbits/sec
      [4]:   100000 kbits/sec, [5]:   100000 kbits/sec
      [6]:   100000 kbits/sec, [7]:    98000 kbits/sec
    Affinity: 0x00000001
    Admin. Weight: 2000
    Physical LINK BW: 100000 kbits/sec
  IP Address:   10.255.255.2
  Metric: 0          IP 10.255.255.2/32
```

```
    Route Admin Tag: 1
    Prefix-attr: X:0 R:0 N:1
    Source Router ID: 10.255.255.2
  Metric: 1          IP 10.1.2.0/24
    Prefix-attr: X:0 R:0 N:0
  Metric: 10         IP 10.2.5.0/24
    Prefix-attr: X:0 R:0 N:0
  Metric: 5          IP 10.2.6.0/24
    Prefix-attr: X:0 R:0 N:0


CSR-1(config)#do sh mpls traffic-eng tunnels tunnel 10


Name: TE-10.254.10.10-Affinity_0x00000001 (Tunnel10) Destination: 10.255.255.10
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected
    path option 1, type dynamic (Basis for Setup, path weight 6504)


  Config Parameters:
    Bandwidth: 1000     kbps (Global)  Priority: 7  7   Affinity: 0x1/0xFFFF
    Metric Type: TE (default)
    Path-selection Tiebreaker:
      Global: not set   Tunnel Specific: not set   Effective: min-fill (default)
    Hop Limit: disabled
    Cost Limit: disabled
    Path-invalidation timeout: 10000 msec (default), Action: Tear
    AutoRoute: disabled LockDown: disabled Loadshare: 1000 [2000000] bw-based
    auto-bw: disabled
    Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
  Active Path Option Parameters:
    State: dynamic path option 1 is active
    BandwidthOverride: disabled  LockDown: disabled  Verbatim: disabled
  Node Hop Count: 4


  InLabel  :  -
```

```
  OutLabel : GigabitEthernet3, 20046
  Next Hop : 10.1.2.2
  RSVP Signalling Info:
       Src 10.255.255.1, Dst 10.255.255.10, Tun_Id 10, Tun_Instance 1
    RSVP Path Info:
      My Address: 10.1.2.1
      Explicit Route: 10.1.2.2 10.2.5.5 10.5.8.8 10.8.10.10
                      10.255.255.10
      Record   Route:   NONE
      Tspec: ave rate=1000 kbits, burst=1000 bytes, peak rate=1000 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=1000 kbits, burst=1000 bytes, peak rate=1000 kbits
  Shortest Unconstrained Path Info:
    Path Weight: 6503 (TE)
    Explicit Route: 10.1.4.4 10.4.6.6 10.6.9.9 10.9.10.10
                    10.255.255.10

  History:
    Tunnel:
      Time since created: 5 hours, 58 minutes
      Time since path change: 5 minutes, 39 seconds
      Number of LSP IDs (Tun_Instances) used: 1
    Current LSP: [ID: 1]
      Uptime: 5 minutes, 39 seconds

CSR-1(config)#do sh mpls traffic-eng tunnels tunnel 11

Name: TE-10.254.10.10-Affinity_0x00000002 (Tunnel11) Destination: 10.255.255.10
  Status:
    Admin: up         Oper: down   Path: not valid   Signalling: Down
    path option 1, type dynamic

  Config Parameters:
```

```
  Bandwidth: 1000     kbps (Global)  Priority: 7  7   Affinity: 0x2/0xFFFFFFFF
  Metric Type: TE (default)
  Path-selection Tiebreaker:
    Global: not set   Tunnel Specific: not set   Effective: min-fill (default)
  Hop Limit: disabled
  Cost Limit: disabled
  Path-invalidation timeout: 10000 msec (default), Action: Tear
  AutoRoute: disabled LockDown: disabled Loadshare: 1000 [0] bw-based
  auto-bw: disabled
  Fault-OAM: disabled, Wrap-Protection: disabled, Wrap-Capable: No
Shortest Unconstrained Path Info:
  Path Weight: 6503 (TE)
  Explicit Route: 10.1.4.4 10.4.6.6 10.6.9.9 10.9.10.10
                  10.255.255.10


History:
  Tunnel:
    Time since created: 6 hours
    Number of LSP IDs (Tun_Instances) used: 16
```

Task 17:

- Shutdown all interfaces other than links between R1, R3, R5, R6
- The interface between R6 and R8 should also be UP
- The GigabitEthernet1 interface of R3 should be protected using RSVP TE tunnel between R3 and R6 (Tunnel 3366 as being NHOP Backup Tunnel)
- Create Tunnel 9 on R1 and force it to support the FRR link protection that R3 is providing
- The Tunnel 9 destination should be 10.255.255.8

Solution:

NOTE: before you continue doing the task, Shutdown the interfaces that mentioned in the task.

The RSVP TE tunnels in the SP networks are mostly used for FRR.

In this task we will configure Link Protection using RSVP TE tunnels.

- CSR-1 is the Head-End router
- CSR-3 will be the PLR (Point of Local Repair) and it is going to protect it's GigabitEthernet1 interface
- CSR-6 will be the MP (Merge Point)

First of all, a TE Tunnel should be formed between CSR-3 (PLR) and CSR-6 (MP), and an explicit path will be used in order to form this LSP, then we can use this tunnel as a backup tunnel to protect GigabitEthernet1 of CSR-3 (PLR). As soon as this interface fails, the backup tunnel will be used to forward the packets (This switchover happens under 50 ms).

```
CSR-3:
ip explicit-path name R3-R5 enable
 index 1 next-address 10.3.5.5
 index 2 next-address 10.5.6.6
!
interface Tunnel3366
 description RSVP-_LP_PLR_MP_Tunnel
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.6
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
```

```
 tunnel mpls traffic-eng path-option 1 explicit name R3-R5
!
interface GigabitEthernet1
 mpls traffic-eng backup-path tunnel 3366
!
```

CSR-1:

```
interface Tunnel9
 description TE-FRR_Link_Protection
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.8
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng path-selection metric igp
 tunnel mpls traffic-eng fast-reroute
!
ip route 10.253.9.8 255.255.255.255 Tunnel9 name TE-FRR_Link_Protection
```

That is all you need to do!

Verification:

```
CSR-3(config)#
*Apr 15 00:50:48.126: %BFD-6-BFD_SESS_CREATED: BFD-SYSLOG: bfd_session_created, neigh 10.3.6.6 proc:FRR,
idb:GigabitEthernet1 handle:1 act


CSR-3(config)#do sh 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
---------------------------    -------- --------------   --------------   ------
10.255.255.1 9 [3]             30050    Gi1:60025        Tu3366:60025     ready
```

```
CSR-3(config)#do sh mpls traffic-eng fast-reroute  database detail
FRR Database Summary:
  Protected interfaces    : 1
  Protected LSPs/Sub-LSPs : 1
  Backup tunnels          : 1
  Active interfaces       : 0
  FRR Active tunnels      : 0


P2P LSPs:


 Tun ID: 9, LSP ID: 3, Source: 10.255.255.1
 Destination: 10.255.255.8
  State         : ready
  InLabel       : 30050
  OutLabel      : Gi1:60025
  FRR OutLabel  : Tu3366:60025
```

As you can realize, the backup tunnel is ready to forward the packets over this tunnel whenever Gig1 fails (under 50ms), And that packet will have two labels, a label to reach the CSR-5 and another inner label for that failed link: 60025.

Task 18:

- Shutdown all interfaces other than links between R1, R3, R5, R6, R8
- The link between R8 and R10 should be UP
- The GigabitEthernet1 interface of R3 should be protected using RSVP TE tunnel between R3 and R8 (Tunnel 3388 as being NNHOP Backup Tunnel)
- Create Tunnel 12 on R1 and force it to support the FRR Node Protection that R3 is providing (Node Protection for R6)
- The Tunnel 12 destination should be 10.255.255.10

Solution:

NOTE: before you continue doing the task, Shutdown the interfaces that mentioned in the task.

This time, in addition to link protection, we can achieve the NODE Protection with RSVP TE Tunnels.

- CSR-1 will be the Head-End device
- CSR-3 will be the PLR protecting the Gig1 link and also CSR-6 node failure
- CSR-8 will be the Merge Point (MP)

In this scenario if the link between CSR-3 and CSR-6 fails or even CSR-6 node fails completely, the PLR node (CSR-3) could handle this situation (switchover to the Tunnel interface under 50ms).

CSR-3:

```
int tunnel 3366
 shutdown
!
ip explicit-path name R3-R5-R8 enable
 index 1 next-address 10.3.5.5
 index 2 next-address 10.5.8.8
 index 3 next-address 10.255.255.8
!
interface Tunnel3388
 description RSVP-_NP_PLR_MP_Tunnel
 ip unnumbered Loopback0
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.8
```

```
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 explicit name R3-R5-R8
!
interface GigabitEthernet1
 no mpls traffic-eng backup-path Tunnel3366
 mpls traffic-eng backup-path Tunnel3388
!
```

CSR-1:

```
int tunnel 9
 shutdown
!
interface Tunnel12
 description TE-FRR_Node_Protection
 ip unnumbered Loopback0
 no shutdown
 tunnel mode mpls traffic-eng
 tunnel destination 10.255.255.8
 tunnel mpls traffic-eng priority 7 7
 tunnel mpls traffic-eng bandwidth 1000
 tunnel mpls traffic-eng path-option 1 dynamic
 tunnel mpls traffic-eng path-selection metric igp
 tunnel mpls traffic-eng fast-reroute node-protect
!
```

CSR-5:

```
interface GigabitEthernet2
 no  mpls traffic-eng attribute-flags 0x1
!
```

Now, everything seems to be ready.

Let's do the verification:

Verification:

```
CSR-3(config)#

*Apr 15 01:16:25.981: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel3388, changed state to up

CSR-3(config)#

*Apr 15 01:16:26.010: %BFD-6-BFD_SESS_CREATED: BFD-SYSLOG: bfd_session_created, neigh 10.3.6.6 proc:FRR,
idb:GigabitEthernet1 handle:1 act

CSR-3(config)#do sh 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

---------------------------    -------- --------------   --------------   ------

10.255.255.1 12 [10]           30053    Gi1:60046        Tu3388: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.

CSR-3(config)#do sh ip rsvp fast-reroute

P2P                        Protect BW         Backup

Protected LSP              I/F     BPS:Type   Tunnel:Label  State  Level   Type

-------------              ------- --------   ------------ ------ -----   ------

TE-FRR_Node_Protection     Gi1     1M:G       Tu3388:3      Ready  any-unl N-Nhop


P2MP

*Protected Sub-LSP                       Protect BW         Backup

src_lspid[subid]->dst_tunid              I/F     BPS:Type   Tunnel:Label  State

---------------------------              ------- --------   ------------ ------
```