# **MPLS RSVP-TE and VPN LAB 1**

Matt Capranos
Tom Woo
Roland Szczesny
Stuart MacLean

## Executive Summary

The purpose of this lab was to implement MPLS and determine how it is used for network traffic engineering. Effective traffic engineering in a network depends on four components: packet forwarding, information distribution, path selection and signaling. The packet forwarding is provided by MPLS (Multiprotocol Label Switching) which is used for directing packets along a route; this route is known as a LSP (Label Switched Path), which is pre-determined by another protocol, another component that is needed in TE is information distribution.

Information Distribution is used to build a topology of the network which will include characteristics which will be used in selecting the LSR's (Label Switch Routers) that will make up the Label Switched Path. This information is distributed by using Link State type 10 Opaque messages extension in OSPF. OSPF will include link attributes in its advertisements to routers in its domain.

The third component is path selection and it takes place in the ingress Label Switch Router (LSR), all routers receive IGP updates which include the link attributes. The link attributes and topology information are stored in the TED (traffic engineering database) which is used in determining the path that will be used by the LSP.

The path in this lab was a strict explicit route, meaning that the LSR's that where to be used along the path where manually selected in the Ingress LSR. Dynamic path selection occurs when some or none of the LSR's path is specified in the Ingress LSR. This process is known as loose explicit route selection. The dynamic path is selected by applying CSPF (Constrained Shortest Path First) to the TED (which contains the network topology and the attributes of the routers in the domain) in the Ingress LSR. The route, whether it is loose or strict is then passed to the fourth and final component, the signaling process.

The signaling component engages the forwarding state, initializes the LSP state and is responsible for label distribution. Signaling is provided by the RSVP (Resource Reservation Protocol) and its extensions. RSVP uses rsvp_path and rsvp_resv messages, the rsvp_path message will request that the LSR, which will be in the LSP, should provide a label binding for the LSP that RSVP is establishing. The rsvp_resv messages take the reverse path that the rsvp_path messages took and rsvp_resv distributes the labels to the LSR's. So, put simply, rsvp_path gets the LSR ready for labels and rsvp_resv gives out the labels. RSVP is also responsible for rerouting of LSP's and for the signaling of LSP teardowns.

RSVP signaling was studied in detail, the transactions that are dedicated to establishing connections, reservations of resources and maintenance of the label switched paths where analyzed. Table 3s router (3.3.3.3) was used as our ingress LSR router and router 10.10.10.10 was picked as out egress LSR router, router 1.1.1.1 was used for our destination.

A strict explicit path was configured for our MPLS tunnel. The path did not follow Shortest Path First and this was verified by doing traceroutes to router 1.1.1.1 with the tunnel active and again with the tunnel shutdown
.
A policy was set up to map traffic from our network to router 1.1.1.1's network using our tunnel.  Sniffing was used to collect RSVP protocol messages.  The path and resources reservation transactions of RSVP where sniffed and analyzed.  **Tunnel or path maintenance data was also discovered.**  The tunnel was shut down in the Ingress LSR and the tear down process was sniffed and analyzed.

## LSP path for MPLS Tunnel

For our lab the destination for our LSP path was router 1.1.1.1, the explicit path for that route to follow is found below.

```
ip explicit-path name R3_R10nable
 next-address 11.11.11.11
 next-address 7.7.7.7
 next-address 8.8.8.8
 next-address 10.10.10.10
```

The explicit route indicates the path that the VPN tunnel will use to reach our destination, the configuration for our tunnel is found below.

```
interface Tunnel0
 bandwidth 100
 ip unnumbered Loopback0
 load-interval 30
 mpls traffic-eng tunnels
 tunnel destination 1.1.1.1
 tunnel mode mpls traffic-eng
 tunnel mpls traffic-eng autoroute announce
 tunnel mpls traffic-eng priority 1 1
 tunnel mpls traffic-eng bandwidth 100
 tunnel mpls traffic-eng path-option 10 explicit name R3_R10
 tunnel mpls traffic-eng record-route
```

Without the use of a VPN tunnel with explicit route set, the IGP shortest path would be used to reach our destination 1.1.1.1 (during the time of the traceroute, router 2.2.2.2 was not running) , the IGP shortest path is found below.

```
R3#traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 192.168.12.11 4 msec 4 msec 0 msec
  2 10.0.0.1 4 msec 4 msec *
```

After the Tunnel was taken from a state of **shutdown** to **no shutdown**, the VPN tunnel was up and operational again, doing the same traceroute from router 3.3.3.3 to router 1.1.1.1 now uses the VPN tunnel.

```
R3#traceroute 1.1.1.1
Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 192.168.12.11 [MPLS: Label 16 Exp 0] 8 msec 8 msec 4 msec
  2 192.168.11.7 [MPLS: Label 16 Exp 0] 8 msec 8 msec 4 msec
  3 10.0.78.1 [MPLS: Label 16 Exp 0] 4 msec 4 msec 4 msec
  4 192.168.11.10 [MPLS: Label 16 Exp 0] 4 msec 4 msec 4 msec
  5 10.0.101.1 4 msec 4 msec *
```

**Policy to map traffic from our network to another the tunnel**

This step of the lab required a policy to be put in place to map traffic from our network to another using the VPN tunnel that was put in place. For our policy map a simple policy map was created, the example and explanation of the map are found below.

```
class-map match-all tunnelR3_R10map
  match any
!
!
policy-map tunnelR3_R10
!
```

**Class map** is used to specify the class map to which packets will be matched; in our case we are matching all with **match-all** the destination the packets will be sent to will be **tunnelR3_R10map**. The tunnelR3_R10 is selected with the use of **policy-map tunnelR3_R10.**

## Transactions for path/resources reservation

In order to establish a tunnel between router 3.3.3.3 and router 1.1.1.1, there will be communication between these 2 routers. In the sniff, we see that router 3.3.3.3 is sending router 1.1.1.1 path messages about the explicit path of a tunnel going through router 11.11.11.11, router 7.7.7.7, router 8.8.8.8, and then through router 10.10.10.10 to router 1.1.1.1. Due to the size of the path message, the path message is separated into 2 path messages, as illustrated below.

| No. | Time | Source | Destination | Protocol | Info |
|---|---|---|---|---|---|
| 59 | 31.837063 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0 |
| 60 | 31.858587 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0 |

The picture below shows the first part of the path message, illustrating the explicit route the traffic will be going through when router 1.1.1.1 is the destination. The rsvp_path message appeared to have been sent in 2 pieces towards our egress router, the first packet shown below was sent at 31.858587 seconds.

```
⊟ EXPLICIT ROUTE: IPv4 192.168.11.7, IPv4 10.0.78.1, IPv4 192.168.11.8, ...
    Length: 52
    Class number: 20 - EXPLICIT ROUTE object
    C-type: 1
  ⊞ IPv4 Subobject - 192.168.11.7, Strict
  ⊞ IPv4 Subobject - 10.0.78.1, Strict
  ⊞ IPv4 Subobject - 192.168.11.8, Strict
  ⊞ IPv4 Subobject - 192.168.11.10, Strict
  ⊞ IPv4 Subobject - 10.0.101.1, Strict
  ⊞ IPv4 Subobject - 1.1.1.1, Strict
⊟ LABEL REQUEST: Basic: L3PID: IP (0x0800)
```

The second rsvp_path message was sent at (31.899829) almost the exact same time completing the path from our ingress to egress router, as shown below.

```
⊞ HOP: IPv4, 192.168.11.8
⊞ TIME VALUES: 30000 ms
⊟ EXPLICIT ROUTE: IPv4 192.168.11.10, IPv4 10.0.101.1, IPv4 1.1.1.1
    Length: 28
    Class number: 20 - EXPLICIT ROUTE object
    C-type: 1
  ⊞ IPv4 Subobject - 192.168.11.10, Strict
  ⊞ IPv4 Subobject - 10.0.101.1, Strict
  ⊞ IPv4 Subobject - 1.1.1.1, Strict
⊞ LABEL REQUEST: Basic: L3PID: IP (0x0800)
```

After the path messages have been sent, the egress routers reply with resv messages to reserve/maintain the tunnel that has been established as shown below.

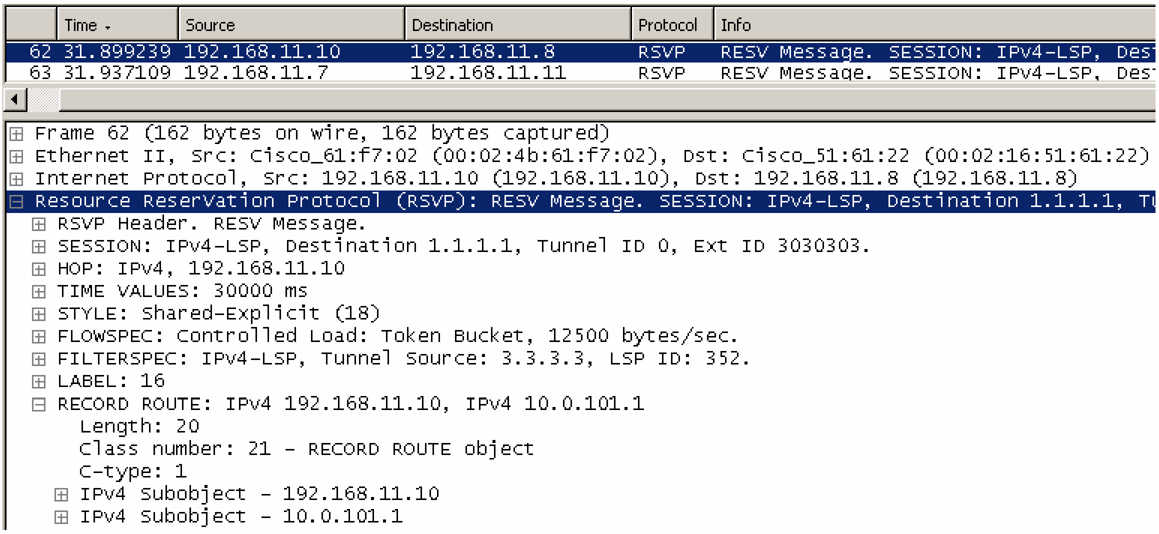


We did not see the resv message between router 8.8.8.8 and router 7.7.7.7, as the path between the 2 routers does not go through the switch where were sniffing the traffic.

## Data indicating tunnel /path maintenance

The data that indicates tunnel / path maintenance would be the periodic resending of the rsvp_path message between the ingress and egress routers. Below shows a picture of the time stamps that the rsvp_path messages were sent.

Filter: rsvp | Expression... | Clear | Apply

| No. | Time | Source | Destination | Protocol | Info |
|---|---|---|---|---|---|
| 95 | 50.808438 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 109 | 55.939096 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 184 | 90.291896 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 191 | 95.500453 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 212 | 109.52076 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 218 | 113.25239 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 243 | 125.26093 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |

The tunnel maintenance rsvp_path messages are identical to the path message sent during the initial tunnel setup, as shown below.

```
⊞ Internet Protocol, Src: 3.3.3.3 (3.3.3.3), Dst: 1.1.1.1 (1.1.1.1)
⊟ Resource ReserVation Protocol (RSVP): PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 3030303. SEND
  ⊞ RSVP Header. PATH Message.
  ⊞ SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 3030303.
  ⊞ HOP: IPv4, 192.168.11.8
  ⊞ TIME VALUES: 30000 ms
  ⊟ EXPLICIT ROUTE: IPv4 192.168.11.10, IPv4 10.0.101.1, IPv4 1.1.1.1
      Length: 28
      Class number: 20 - EXPLICIT ROUTE object
      C-type: 1
    ⊞ IPv4 Subobject - 192.168.11.10, Strict
    ⊞ IPv4 Subobject - 10.0.101.1, Strict
    ⊞ IPv4 Subobject - 1.1.1.1, Strict
  ⊞ LABEL REQUEST: Basic: L3PID: IP (0x0800)
  ⊞ SESSION ATTRIBUTE: SetupPrio 1, HoldPrio 1, SE Style,  [R3_t0]
  ⊞ SENDER TEMPLATE: IPv4-LSP, Tunnel Source: 3.3.3.3, LSP ID: 352.
  ⊞ SENDER TSPEC: IntServ: Token Bucket, 12500 bytes/sec.
  ⊞ ADSPEC
  ⊞ RECORD ROUTE: IPv4 192.168.11.8, IPv4 10.0.78.2, IPv4 192.168.11.11,
```

In addition to periodic rsvp_path messages being sent to the egress router, rsvp_resv messages were sent from the egress router to the ingress router.

| No. | Time | Source | Destination | Protocol | Info |
|---|---|---|---|---|---|
| 63 | 31.937109 | 192.168.1 | 192.168.1 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 132 | 66.640397 | 192.168.1 | 192.168.1 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 134 | 67.412824 | 192.168.1 | 192.168.1 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 135 | 67.428824 | 192.168.1 | 192.168.1 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 194 | 97.972452 | 192.168.1 | 192.168.1 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 210 | 107.97672 | 192.168.1 | 192.168.1 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 235 | 122.22447 | 192.168.1 | 192.168.1 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |

Again the rsvp_resv messages are identical to the messages being sent during the initial tunnel setup.

One additional note that may help to show tunnel / path maintenance would be the time stamps that the rsvp_path and rsvp_resv messages are sent. During the initial tunnel setup the rsvp_path messages are sent and the rsvp_resv reply message is received at almost identical time stamps.

| No. | Time | Source | Destination | Protocol | Info |
|---|---|---|---|---|---|
| 59 | 31.837063 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 60 | 31.858587 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 62 | 31.899239 | 192.168.1: | 192.168.1: | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 63 | 31.937109 | 192.168.1: | 192.168.1: | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |

In the subsequent rsvp_path and rsvp_resv messages that are sent the time stamps are off by several seconds.

rsvp_path messages

| No. | Time | Source | Destination | Protocol | Info |
|---|---|---|---|---|---|
| 95 | 50.808438 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |
| 109 | 55.939096 | 3.3.3.3 | 1.1.1.1 | RSVP | PATH Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 30303C |

rsvp_resv messages

| No. | Time | Source | Destination | Protocol | Info |
|---|---|---|---|---|---|
| 132 | 66.640397 | 192.168.11.10 | 192.168.11.8 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext |
| 134 | 67.412824 | 192.168.11.10 | 192.168.11.8 | RSVP | RESV Message. SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext |

## Descriptions of tunnel tear down

RSVP Tear down procedure:

When a reserved path is put into a state where the reserved path needs to be torn down, several steps go into effect to remove the reserved path and keeping the network up and operational.

If the sender terminates the connection, a PATH TEAR message is sent. (When the receiver terminates it will send a RESV TEAR message). When a PATH TEAR arrives at a router, it removes its PATH State. Subsequently the RESV state is removed (because when no matching PATH state is in place the RESV state has no meaning of existence). The PATH TEAR is forwarded to the next machine in the tunnel path. RESV TEAR messages will be sent to the previous-hop when a RESV state is removed.

```
⊟ RSVP Header. RESV TEAR Message.
     RSVP Version: 1
     Flags: 00
     Message Type: RESV TEAR Message.  (6)
     Message Checksum: 0xfa2a [correct]
     Sending TTL: 255
     Message length: 100
⊟ SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 3030303.
     Length: 16
     Class number: 1 - SESSION object
     C-type: 7 - IPv4 LSP
     Destination address: 1.1.1.1 (1.1.1.1)
     Tunnel ID: 0
     Extended Tunnel ID: 50529027 (3.3.3.3)
⊟ HOP: IPv4, 192.168.11.7
     Length: 12
     Class number: 3 - HOP object
     C-type: 1 - IPv4
     Neighbor address: 192.168.11.7
     Logical interface: 671089671
⊞ STYLE: Shared-Explicit (18)
⊞ FLOWSPEC: Controlled Load: Token Bucket, 12500 bytes/sec.
⊞ FILTERSPEC: IPv4-LSP, Tunnel Source: 3.3.3.3, LSP ID: 352.
⊞ CONFIRM: Receiver 192.168.11.7
```

This is the first packet in the tear down process. The packet contains the RSVP Header containing the RESV TEAR MESSAGE.

This message is being sent to 1.1.1.1 which is the last router in the reserved path.

```
⊟ Resource Reservation Protocol (RSVP): RESV TEAR CONFIRM Message. SESSION: IPv4-LSP, Destination 1.1.1.1
  ⊟ RSVP Header. RESV TEAR CONFIRM Message.
       RSVP Version: 1
       Flags: 00
       Message Type: RESV TEAR CONFIRM Message.  (10)
       Message Checksum: 0x232a [correct]
       Sending TTL: 255
       Message length: 100
  ⊟ SESSION: IPv4-LSP, Destination 1.1.1.1, Tunnel ID 0, Ext ID 3030303.
       Length: 16
       Class number: 1 - SESSION object
       C-type: 7 - IPv4 LSP
       Destination address: 1.1.1.1 (1.1.1.1)
       Tunnel ID: 0
       Extended Tunnel ID: 50529027 (3.3.3.3)
```

This is the RESV TEAR CONFIRM packet that is sent back to confirm that the reserved path has been removed and that the PATH TEAR message has been successfully forwarded down the chain to reach the final destinations.

All routers must receive a PATH TEAR packet to take down the reserved path, and all sending routers must receive a RESV TEAR message from their destination to confirm that the reserved path has been taken down properly.