Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

Paul Negron
CCIE #14856
SP

VOL-III

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Table of Content:

Subject	Page	Volume
LDP		
Lab 1 Configuring Label Distribution Protocol	5	Vol-III
VRF-LITE		
Lab 1 PBR Based on Source IP Address	75	Vol-III
MPLS & L3VPNs		
Lab 1 Static & RIPv2 Routing in a VPN	82	Vol-III
Lab 2 OSPF Routing in a VPN	111	Vol-III
Lab 3 Backdoor links & OSPF	131	Vol-III
Lab 4 EIGRP Routing in a VPN	147	Vol-III
Lab 5 EIGRP SOO	162	Vol-III
Lab 5 BGP Routing in a VPN	176	Vol-III
Lab 6 Complex VPNs and Filters	193	Vol-III
MPLS and NAT		
Lab 1 MPLS and NAT – I	212	Vol-III
Lab 2 MPLS and NAT – II	225	
Carrier Supporting Carrier		
Lab 1 Carrier Supporting Carrier	235	Vol-III
Lab 2 Inter-AS VPN	248	
VPN Extranet & Internet Access		
Lab 1 Internet Access Methods Using Separate VPN & Dual PVC	261	Vol-III
Lab 2 Static Internet Access Using Global Next-hop	276	Vol-III
Multicast VPN (mVPN)		
Lab 1 Multicast VPN (mVPN)	283	Vol-III
Layer 3 VPNs 6PE/6VPE		
Lab 1 6PE	300	Vol-III
Lab 2 6VPE	312	Vol-III
Layer 2 VPNs		
Lab 1 Confguring MPLS & layer 2 VPNs (AToM)	324	Vol-III
Lab 2 L2TPv3 VPN	335	Vol-III
Lab 3 L2TPv3 Internetworking	338	Vol-III
MPLS Traffic Engineering		
Lab 1 Confguring Static Tunnels	348	Vol-III
Lab 2 Dynamic Tunnels	365	Vol-III
Lab 3 Load Balancing	379	Vol-III
Lab 4 Fast Reroute backup Tunnels	390	Vol-III
Lab 5 MPLS QoS	401	Vol-III

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IOS-XR		
Lab 1 IOS-XR Quick Reference	416	Vol-III
IOS-XR and Layer 3 VPNs		
Lab 1 Static and RIPv2 Routing in a VPN	426	Vol-III
Lab 2 OSPF Routing in a VPN	459	Vol-III
Lab 3 Backdoor links and OSPF	478	Vol-III
Lab 4 Eigrp Routing in a VPN	496	Vol-III
Lab 5 BGP Routing ina VPN	513	Vol-III
IOS-XR Traffic Engineering		
Lab 1 Configuring Static Tunnels	526	Vol-III
Lab 2 Dynamic Tunnels	544	Vol-III
VPWS and VPLS		
Lab 1 VPWS	556	Vol-III
Lab 2 VPLS (EP-LAN)	565	Vol-III

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

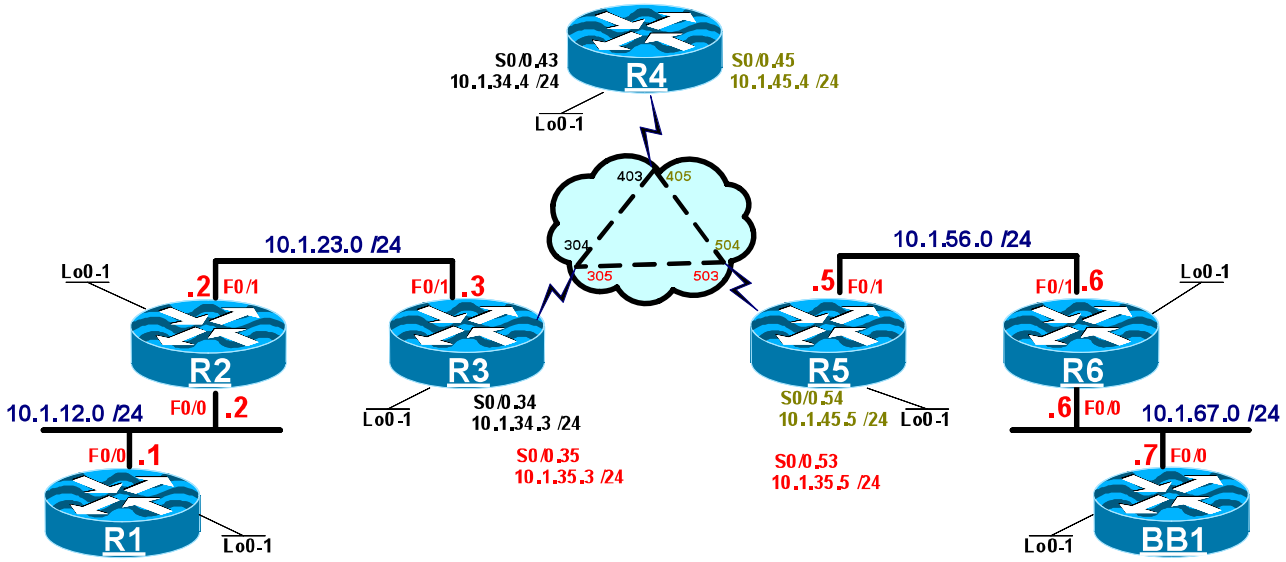
Paul Negron
CCIE #14856
SP

LDP

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1
Configuring Label Distribution Protocol



Lab Setup:

- The frame-relay connections should be configured in a Point-to-Point manner
- The F0/0 interface of R1 and R2 should be configured in VLAN 100
- The F0/1 interface of R2 and R3 should be configured in VLAN 200
- The F0/1 interface of R5 and R6 should be configured in VLAN 300
- The F0/0 interface of R6 and BB1 should be configured in VLAN 400
- Use the IP addressing identified in the following chart.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing chart:

Router	Interface / IP Addressing	Connecting To:
R1	F0/0 = 10.1.12.1 /24 Lo0 = 1.1.1.1 /24 Lo1 = 11.1.1.1 /24	R2's F0/0
R2	F0/0 = 10.1.12.2 /24 F0/1 = 10.1.23.2 /24 Lo0 = 2.2.2.2 /24 Lo1 = 22.2.2.2 /24	R1's F0/0 R3's F0/1
R3	F0/1 = 10.1.23.3 /24 S0/0.34 = 10.1.34.3 /24 S0/0.35 = 10.1.35.3 /24 Lo0 = 3.3.3.3 /24 Lo1 = 33.3.3.3 /24	R2's F0/1 R4's S0/0.43 R5's S0/0.53
R4	S0/0.43 = 10.1.34.4 /24 S0/0.45 = 10.1.45.4 /24 Lo0 = 4.4.4.4 /24 Lo1 = 44.4.4.4 /24	R3's S0/034 R5's S0/0.54
R5	S0/0.53 = 10.1.35.5 /24 S0/0.54 = 10.1.45.5 /24 F0/1 = 10.1.56.5 /24 Lo0 = 5.5.5.5 /24 Lo1 = 55.5.5.5 /24	R3's S0/0.35 R4's S0/0.45 R6's F0/1
R6	F0/1 = 10.1.56.6 /24 F0/0 = 10.1.67.6 /24 Lo0 = 6.6.6.6 /24 Lo1 = 66.6.6.6 /24	R5's F0/1 BB1's F0/0
BB1	F0/0 = 10.1.67.7 /24 Lo0 = 7.7.7.7 /24 Lo1 = 77.7.7.7 /24	R6's F0/0

Task 1

Configure OSPF area 0 on every router and their directly connected interface/s except Loopback1. Ensure that lo0 interface of these router are advertised with their correct mask.

On R1

```
R1(config)#Router ospf 1
R1(config-router)#Netw 10.1.12.1 0.0.0.0 area 0
R1(config-router)#Netw 1.1.1.1 0.0.0.0 area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#Int lo0
R1(config-if)#ip ospf netw point-to-point
```

On R2

```
R2(config)#Router ospf 1
R2(config-router)#Netw 10.1.12.2 0.0.0.0 are 0
R2(config-router)#Netw 10.1.23.2 0.0.0.0 are 0
R2(config-router)#Netw 2.2.2.2 0.0.0.0 are 0

R2(config-router)#int lo0
R2(config-if)#ip ospf netw point-to-point
```

On R3

```
R3(config)#Router ospf 1
R3(config-router)#Netw 10.1.23.3 0.0.0.0 are 0
R3(config-router)#Netw 10.1.34.3 0.0.0.0 are 0
R3(config-router)#Netw 10.1.35.3 0.0.0.0 are 0
R3(config-router)#Netw 3.3.3.3 0.0.0.0 are 0

R3(config-router)#int lo0
R3(config-if)#ip ospf netw point-to-point
```

On R4

```
R4(config)#Router ospf 1
R4(config-router)#Netw 10.1.34.4 0.0.0.0 are 0
R4(config-router)#Netw 10.1.45.4 0.0.0.0 are 0
R4(config-router)#Netw 4.4.4.4 0.0.0.0 are 0

R4(config-router)#int lo0
R4(config-if)#ip ospf netw point-to-point
```

On R5

```
R5(config)#Router ospf 1
R5(config-router)#Netw 10.1.45.5 0.0.0.0 are 0
R5(config-router)#Netw 10.1.56.5 0.0.0.0 are 0
R5(config-router)#Netw 10.1.35.5 0.0.0.0 are 0
R5(config-router)#Netw 5.5.5.5 0.0.0.0 are 0

R5(config-router)#int lo0
R5(config-if)#ip ospf netw point-to-point
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R6

```
R6(config)#Router ospf 1
R6(config-router)#Netw 10.1.56.6 0.0.0.0 are 0
R6(config-router)#Netw 10.1.67.6 0.0.0.0 are 0
R6(config-router)#Netw 6.6.6.6 0.0.0.0 are 0

R6(config-router)#int lo0
R6(config-if)#ip ospf netw point-to-point
```

On BB1

```
BB1(config)#Router ospf 1
BB1(config-router)#Netw 10.1.67.7 0.0.0.0 are 0
BB1(config-router)#Netw 7.7.7.7 0.0.0.0 are 0

BB1(config-router)#int lo0
BB1(config-if)#ip ospf netw point-to-point
```

To verify the configuration:

On R1

```
R1#Show ip route ospf | Inc O

O       2.2.2.0 [110/2] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       3.3.3.0 [110/3] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       4.4.4.0 [110/67] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       5.5.5.0 [110/67] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       6.6.6.0 [110/68] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       7.7.7.0 [110/69] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       10.1.23.0 [110/2] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       10.1.45.0 [110/130] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       10.1.35.0 [110/66] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       10.1.34.0 [110/66] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       10.1.56.0 [110/67] via 10.1.12.2, 00:00:09, FastEthernet0/0
O       10.1.67.0 [110/68] via 10.1.12.2, 00:00:09, FastEthernet0/0
```

On R2

```
R2#Show ip route ospf | Inc O

O       1.1.1.0 [110/2] via 10.1.12.1, 00:02:34, FastEthernet0/0
O       3.3.3.0 [110/2] via 10.1.23.3, 00:02:34, FastEthernet0/1
O       4.4.4.0 [110/66] via 10.1.23.3, 00:02:34, FastEthernet0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       5.5.5.0 [110/66] via 10.1.23.3, 00:02:34, FastEthernet0/1
O       6.6.6.0 [110/67] via 10.1.23.3, 00:02:34, FastEthernet0/1
O       7.7.7.0 [110/68] via 10.1.23.3, 00:02:34, FastEthernet0/1
O       10.1.45.0 [110/129] via 10.1.23.3, 00:02:34, FastEthernet0/1
O       10.1.35.0 [110/65] via 10.1.23.3, 00:02:34, FastEthernet0/1
O       10.1.34.0 [110/65] via 10.1.23.3, 00:02:34, FastEthernet0/1
O       10.1.56.0 [110/66] via 10.1.23.3, 00:02:34, FastEthernet0/1
O       10.1.67.0 [110/67] via 10.1.23.3, 00:02:34, FastEthernet0/1
```

On R3

```
R3#Show ip route ospf | Inc O

O       1.1.1.0 [110/3] via 10.1.23.2, 00:03:18, FastEthernet0/1
O       2.2.2.0 [110/2] via 10.1.23.2, 00:03:18, FastEthernet0/1
O       4.4.4.0 [110/65] via 10.1.34.4, 00:03:18, Serial0/0.34
O       5.5.5.0 [110/65] via 10.1.35.5, 00:03:18, Serial0/0.35
O       6.6.6.0 [110/66] via 10.1.35.5, 00:03:18, Serial0/0.35
O       7.7.7.0 [110/67] via 10.1.35.5, 00:03:18, Serial0/0.35
O       10.1.12.0 [110/2] via 10.1.23.2, 00:03:18, FastEthernet0/1
O       10.1.45.0 [110/128] via 10.1.34.4, 00:03:18, Serial0/0.34
O       10.1.56.0 [110/65] via 10.1.35.5, 00:03:18, Serial0/0.35
O       10.1.67.0 [110/66] via 10.1.35.5, 00:03:18, Serial0/0.35
```

On R4

```
R4#Show ip route ospf | Inc O

O       1.1.1.0 [110/67] via 10.1.34.3, 00:04:10, Serial0/0.43
O       2.2.2.0 [110/66] via 10.1.34.3, 00:04:10, Serial0/0.43
O       3.3.3.0 [110/65] via 10.1.34.3, 00:04:10, Serial0/0.43
O       5.5.5.0 [110/65] via 10.1.45.5, 00:04:10, Serial0/0.45
O       6.6.6.0 [110/66] via 10.1.45.5, 00:04:10, Serial0/0.45
O       7.7.7.0 [110/67] via 10.1.45.5, 00:04:10, Serial0/0.45
O       10.1.12.0 [110/66] via 10.1.34.3, 00:04:10, Serial0/0.43
O       10.1.23.0 [110/65] via 10.1.34.3, 00:04:10, Serial0/0.43
O       10.1.35.0 [110/128] via 10.1.34.3, 00:04:10, Serial0/0.43
O       10.1.56.0 [110/65] via 10.1.45.5, 00:04:10, Serial0/0.45
O       10.1.67.0 [110/66] via 10.1.45.5, 00:04:10, Serial0/0.45
```

On R5

```
R5#Show ip route ospf | Inc O

O       1.1.1.0 [110/67] via 10.1.35.3, 00:04:47, Serial0/0.53
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       2.2.2.0 [110/66] via 10.1.35.3, 00:04:47, Serial0/0.53
O       3.3.3.0 [110/65] via 10.1.35.3, 00:04:47, Serial0/0.53
O       4.4.4.0 [110/65] via 10.1.45.4, 00:04:47, Serial0/0.54
O       6.6.6.0 [110/2] via 10.1.56.6, 00:04:47, FastEthernet0/1
O       7.7.7.0 [110/3] via 10.1.56.6, 00:04:47, FastEthernet0/1
O       10.1.12.0 [110/66] via 10.1.35.3, 00:04:47, Serial0/0.53
O       10.1.23.0 [110/65] via 10.1.35.3, 00:04:47, Serial0/0.53
O       10.1.34.0 [110/128] via 10.1.45.4, 00:04:47, Serial0/0.54
O       10.1.67.0 [110/2] via 10.1.56.6, 00:04:47, FastEthernet0/1
```

On R6

```
R6#Show ip route ospf | Inc O

O       1.1.1.0 [110/68] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       2.2.2.0 [110/67] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       3.3.3.0 [110/66] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       4.4.4.0 [110/66] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       5.5.5.0 [110/2] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       7.7.7.0 [110/2] via 10.1.67.7, 00:05:24, FastEthernet0/0
O       10.1.12.0 [110/67] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       10.1.23.0 [110/66] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       10.1.45.0 [110/65] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       10.1.35.0 [110/65] via 10.1.56.5, 00:05:24, FastEthernet0/1
O       10.1.34.0 [110/129] via 10.1.56.5, 00:05:24, FastEthernet0/1
```

On BB1

```
BB1#Show ip route ospf | Inc O

O       1.1.1.0 [110/69] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       2.2.2.0 [110/68] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       3.3.3.0 [110/67] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       4.4.4.0 [110/67] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       5.5.5.0 [110/3] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       6.6.6.0 [110/2] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       10.1.12.0 [110/68] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       10.1.23.0 [110/67] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       10.1.45.0 [110/66] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       10.1.35.0 [110/66] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       10.1.34.0 [110/130] via 10.1.67.6, 00:05:51, FastEthernet0/0
O       10.1.56.0 [110/2] via 10.1.67.6, 00:05:51, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure Label Distribution Protocol on the Ethernet interfaces and Frame-relay subinterfaces interconnecting the routers in this topology. Ensure that the LDP-ID is based on the IP address assigned to the Loopback 0 interface of these routers. You may override a command from the previous task to accomplish this task.

When configuring the label distribution protocol, the actual protocol that creates and distributes labels must be specified, this protocol can be either LDP or TDP; TDP is Cisco proprietary tag distribution protocol which is no longer in use; This lab strickly focuses on LDP and NOT TDP.

The "mpls label protocol" command can be configured to specify the label distribution protocol; this command can be placed under each interface to override the cisco default label protocol, in IOS releases prior to 12.4 TDP is the default protocol used.

When the "mpls label protocol LDP" command is configured in global configuration mode, all interfaces will inherit LDP as the label distribution protocol. This command can also be configured in the interface configuration mode, in which case the label distribution protocol through that interface is affected ONLY.

Each LSR that is running LDP will be assigned a router-id, the "mpls ldp router-id" command can be used to set the router id of a given LSR, the LDP-ID defaults to similar rules as OSPF, Eigrp or BGP, which means that if it is not statically configured, then, numerically highest IP address of any loopback interface is selected as the LDP-ID, if one does not exist, then, the highest active IP address configured on the router is chosen as the LDP-ID.

However, it's good practice to configure it manually to ensure that the transport address of the mpls peer is stable. Remember that LDP advertises its LDP Router-id as the transport address in the LDP discovery hello messages sent from the interface. Therefore, You must provide reachability for that router-id, there must be an exact match for the LDP-ID in the routing table.

The "mpls ip" command enables MPLS forwarding of IPv4 packets along normally routed paths; in some documents this is called Dynamic Label Switching.

This can be enabled globally or on a given interface/s.

On All routers

The following command configures LDP as the label distribution protocol:

```
Rx(config)#mpls label protocol ldp
```

Configuring the LDP-ID of the LSRs based on the IP address of the Lo0 interface:

```
Rx(config)#mpls ldp router-id loopback 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Finally, enabling the MPLS forwarding of IPv4 packets along the normal routed paths:

On R1

```
R1(config)#int F0/0
R1(config-if)#mpls ip
```

On R2

```
R2(config)#int F0/0
R2(config-if)#mpls ip
```

You should get the following console message stating that the local router has discovered a neighbor, in this case R1 (1.1.1.1):

%LDP-5-NBRCHG: LDP Neighbor 1.1.1.1:0 is UP

Explanation of 1.1.1.1:0
Note, LDP uses a 6 Byte quantity, the first 4 Bytes are the LDP-ID or the LDP router-id specified by the "MPLS ldp router-id" command and the last two bytes ":0" identifies the label space, for platform wide label space the last two bytes are always set to "0".

```
R2(config)#int F0/1
R2(config-if)#mpls ip
```

To verify the configuration:

On R2

```
R2#Show mpls interface

Interface              IP            Tunnel   Operational
FastEthernet0/0        Yes (ldp)     No       Yes
FastEthernet0/1        Yes (ldp)     No       Yes
```

Note the output of the above command reveals that LDP is running on F0/0 and F0/1 interfaces of R2.

Let's configure R3:

On R3

```
R3(config)#int F0/1
R3(config-subif)#mpls ip
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-subif)#int S0/0.34
R3(config-subif)#mpls ip

R3(config-subif)#int S0/0.35
R3(config-subif)#mpls ip
```

To verify the configuration:

On R3

```
R3#Show mpls interface

Interface              IP            Tunnel   Operational
FastEthernet0/1        Yes (ldp)     No       Yes
Serial0/0.34           Yes (ldp)     No       Yes
Serial0/0.35           Yes (ldp)     No       Yes
```

Let's configure R4:

On R4

```
R4(config)#int S0/0.43
R4(config-subif)#mpls ip

R4(config-subif)#int S0/0.45
R4(config-subif)#mpls ip
```

To verify the configuration:

On R4

```
R4#Show mpls interfaces

Interface              IP            Tunnel   Operational
Serial0/0.43           Yes (ldp)     No       Yes
Serial0/0.45           Yes (ldp)     No       Yes
```

Let's configure R5:

On R5

```
R5(config)#int F0/1
R5(config-subif)#mpls ip
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-subif)#int S0/0.54
R5(config-subif)#mpls ip

R5(config-subif)#int S0/0.53
R5(config-subif)#mpls ip
```

To verify the configuration:

On R5

```
R5#Show mpls interfaces

Interface              IP             Tunnel   Operational
FastEthernet0/1        Yes (ldp)      No       Yes
Serial0/0.53           Yes (ldp)      No       Yes
Serial0/0.54           Yes (ldp)      No       Yes
```

Let's configure R6:

On R6

```
R6(config)#int F0/1
R6(config-subif)#mpls ip

R6(config)#int F0/0
R6(config-subif)#mpls ip
```

To verify the configuration:

On R6

```
R6#Show mpls interface

Interface              IP             Tunnel   Operational
FastEthernet0/0        Yes (ldp)      No       Yes
FastEthernet0/1        Yes (ldp)      No       Yes
```

Let's configure BB1:

On BB1

```
BB1(config)#int F0/0
BB1(config-if)#mpls ip
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
BB1#Show mpls interfaces

Interface              IP            Tunnel   Operational
FastEthernet0/0        Yes (ldp)     No       Yes
```

To verify the configuration:

On R1

Once the LSRs are configured, they will attempt to discover neighbor/s, this discovery uses hello messages, these messages are UDP packets using 224.0.0.2 port 646; every hello message has a hold timer, by default, the hello messages are sent every 5 seconds, and the hold timer is set to 15 seconds.

To see the discover messages:

```
R1(config)#access-list 100 permit udp host 10.1.12.1 eq 646 any

R1#Debug ip packet det 100
```

IP: s=10.1.12.1 (local), d=224.0.0.2 (FastEthernet0/0), len 62, sending broad/multicast UDP src=646, dst=646

Note the output of the above debug command reveals that the hello messages are sent to destination address of 224.0.0.2 using UDP port 646 as source and destination.

To turn off the debug

On R1

```
R1#U all

R1#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 1.1.1.1:0
        TCP connection: 2.2.2.2.42851 - 1.1.1.1.646
        State: Oper; Msgs sent/rcvd: 52/52; Downstream
        Up time: 00:29:55
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.12.2
        Addresses bound to peer LDP Ident:
          10.1.12.2       10.1.23.2       2.2.2.2         22.2.2.2
```

Note the peer LDP Ident: identifies the LDP-ID of the peer, in this case the peer's LDP-ID is 2.2.2.2:0, the two bytes value of ":0" identifies the Label space; this can be platform-wide or per-interface label space, if the value is zero, it's platform-wide, and anything other than zero is per-interface. In per-

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

interface label space, the packets are forwarded based on both incoming interface and the label, whereas, in platform-wide label space, the LSR generates a label for a given destination and advertises that same label to all of its peers, and the packets are forwarded purley based on the actual label; in Frame mode MPLS, this will always be platform-wide with a value of 0.

Local LDP Ident - identifies the local LSR's LDP-ID (1.1.1.1:0).

In the second line, the TCP connection information is revealed, Note the remote router uses a high port value of 42851 and connects to R1 (1.1.1.1) port 646. Note the high port can be a different value on your router.

The third and the forth lines display the number of messages sent and received and the amount of time that the two LSRs have been up.

The fifth and the sixth lines (LDP Discovery Sources :) display the interface through which the neighbor/s were discovered and it also includes the IP address of the neighbor.

The last line (Addresses bound to peer LDP Ident :) displays the IP addresses that are directly connected to the neighboring LSR.

To verify neighbor 1.1.1.1 on R2:

On R2

```
R2#Show mpls ldp neighbor 1.1.1.1

    Peer LDP Ident: 1.1.1.1:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 1.1.1.1.646 - 2.2.2.2.42851
        State: Oper; Msgs sent/rcvd: 59/60; Downstream
        Up time: 00:36:22
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.12.1
        Addresses bound to peer LDP Ident:
          10.1.12.1       1.1.1.1         11.1.1.1
```

On R3

```
R3#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.51742
        State: Oper; Msgs sent/rcvd: 54/54; Downstream
        Up time: 00:32:42
        LDP discovery sources:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        FastEthernet0/1, Src IP addr: 10.1.23.2
      Addresses bound to peer LDP Ident:
        10.1.12.2       10.1.23.2       2.2.2.2         22.2.2.2
  Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 3.3.3.3:0
      TCP connection: 4.4.4.4.50746 - 3.3.3.3.646
      State: Oper; Msgs sent/rcvd: 52/51; Downstream
      Up time: 00:29:58
      LDP discovery sources:
        Serial0/0.34, Src IP addr: 10.1.34.4
      Addresses bound to peer LDP Ident:
        10.1.34.4       10.1.45.4       4.4.4.4         44.4.4.4
  Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 3.3.3.3:0
      TCP connection: 5.5.5.5.20504 - 3.3.3.3.646
      State: Oper; Msgs sent/rcvd: 50/50; Downstream
      Up time: 00:28:27
      LDP discovery sources:
        Serial0/0.35, Src IP addr: 10.1.35.5
      Addresses bound to peer LDP Ident:
        10.1.35.5       10.1.45.5       10.1.56.5       5.5.5.5
        55.5.5.5
```

On R4

```
R4#Show mpls ldp neighbor

    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 3.3.3.3.646 - 4.4.4.4.50746
        State: Oper; Msgs sent/rcvd: 52/54; Downstream
        Up time: 00:30:57
        LDP discovery sources:
          Serial0/0.43, Src IP addr: 10.1.34.3
        Addresses bound to peer LDP Ident:
          10.1.34.3       10.1.35.3       10.1.23.3       3.3.3.3
          33.3.3.3
    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 5.5.5.5.32239 - 4.4.4.4.646
        State: Oper; Msgs sent/rcvd: 51/51; Downstream
        Up time: 00:29:30
        LDP discovery sources:
          Serial0/0.45, Src IP addr: 10.1.45.5
        Addresses bound to peer LDP Ident:

          10.1.35.5       10.1.45.5       10.1.56.5       5.5.5.5
          55.5.5.5
```

On R5

```
R5#Show mpls ldp neighbor
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 5.5.5.5:0
        TCP connection: 4.4.4.4.646 - 5.5.5.5.32239
        State: Oper; Msgs sent/rcvd: 52/52; Downstream
        Up time: 00:30:02
        LDP discovery sources:
          Serial0/0.54, Src IP addr: 10.1.45.4
        Addresses bound to peer LDP Ident:
          10.1.34.4       10.1.45.4       4.4.4.4         44.4.4.4
    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 5.5.5.5:0
        TCP connection: 3.3.3.3.646 - 5.5.5.5.20504
        State: Oper; Msgs sent/rcvd: 52/52; Downstream
        Up time: 00:29:59
        LDP discovery sources:
          Serial0/0.53, Src IP addr: 10.1.35.3
        Addresses bound to peer LDP Ident:
          10.1.34.3       10.1.35.3       10.1.23.3       3.3.3.3
          33.3.3.3
    Peer LDP Ident: 6.6.6.6:0; Local LDP Ident 5.5.5.5:0
        TCP connection: 6.6.6.6.31410 - 5.5.5.5.646
        State: Oper; Msgs sent/rcvd: 50/50; Downstream
        Up time: 00:28:35
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.56.6
        Addresses bound to peer LDP Ident:
          10.1.67.6       10.1.56.6       6.6.6.6         66.6.6.6
```

On R6

```
R6#Show mpls ldp neighbor

    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 6.6.6.6:0
        TCP connection: 5.5.5.5.646 - 6.6.6.6.31410
        State: Oper; Msgs sent/rcvd: 51/51; Downstream
        Up time: 00:29:19
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.56.5
        Addresses bound to peer LDP Ident:
          10.1.35.5       10.1.45.5       10.1.56.5       5.5.5.5
          55.5.5.5
    Peer LDP Ident: 7.7.7.7:0; Local LDP Ident 6.6.6.6:0
        TCP connection: 7.7.7.7.21915 - 6.6.6.6.646
        State: Oper; Msgs sent/rcvd: 49/50; Downstream
        Up time: 00:27:41
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.67.7
        Addresses bound to peer LDP Ident:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
          10.1.67.7       7.7.7.7          77.7.7.7
```

On BB1

```
BB1#Show mpls ldp neighbor

    Peer LDP Ident: 6.6.6.6:0; Local LDP Ident 7.7.7.7:0
        TCP connection: 6.6.6.6.646 - 7.7.7.7.21915
        State: Oper; Msgs sent/rcvd: 51/50; Downstream
        Up time: 00:28:27
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.67.6
        Addresses bound to peer LDP Ident:
          10.1.67.6       10.1.56.6       6.6.6.6         66.6.6.6
```

To see the discovery hellos and hold time:

```
BB1#Show mpls ldp discovery det

Local LDP Identifier:
    7.7.7.7:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 7.7.7.7
            LDP Id: 6.6.6.6:0; no host route to transport addr
              Src IP addr: 10.1.67.6; Transport IP addr: 6.6.6.6
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Reachable via 6.6.6.0/24
```

Note, there are two types of discovery:

1. **Basic - This type is used to discover directly connected LDP LSRs, these messages are sent to all routers on this subnet out of each interface that has LDP enabled.**

2. **Extended – This type is used between non-directly connected LDP LSRs. For this, an LSR sends targeted hello messages to a specified IP address. This will be described in details in the Traffic engineering section.**

Once the LSRs discover one another, they will attempt to establish an LDP session between them, this session is based on TCP and uses port 646, the TCP connection to an mpls enabled peer will be maintained by the peer with the highest router-id.

Since the router-ids are based on the loopback 0 interface of the routers, this means that R2 will maintain the session.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

In order for this session to establish, the local LSR must have an exact match to the IP address of the neighboring router-id.

Note in the seventh line of the above output the "no host route to transport addr", reveals that the local LSR does NOT have the exact match for the router-id of the peer in its routing table.

The following Show command reveals the same fact:

```
BB1#Show mpls ldp discovery all

Local LDP Identifier:
    7.7.7.7:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 6.6.6.6:0; no host route
```

Note the following command reveals that the network 6.6.6.0/24 is reachable by the local LSR, but NOT the host route 6.6.6.6/32 which is the exact match to the router-id of the neighboring LSR.

```
BB1#Show ip route 6.6.6.0 255.255.255.0

Routing entry for 6.6.6.0/24
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 10.1.67.6 on FastEthernet0/0, 01:35:03 ago
  Routing Descriptor Blocks:
  * 10.1.67.6, from 66.6.6.6, 01:35:03 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1

BB1#Ping 6.6.6.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

BB1#Show ip route 6.6.6.6 255.255.255.255

% Subnet not in table
```

In order to fix this problem, we should remove the "IP OSPF netw point-to-point" interface command from the previous task which advertised the loopback interfaces of the routers with their correct mask.

This will advertise prefix 6.6.6.6/32

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On All routers

```
Rx(config)#int lo0
Rx(config-if)#NO ip ospf netw point-to-point
```

To verify the configuration:

On BB1

```
R1#Show ip route 6.6.6.6 255.255.255.255

Routing entry for 6.6.6.6/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 10.1.67.6 on FastEthernet0/0, 00:02:57 ago
  Routing Descriptor Blocks:
  * 10.1.67.6, from 66.6.6.6, 00:02:57 ago, via FastEthernet0/0
      Route metric is 2, traffic share count is 1

BB1#Show ip route 6.6.6.0 255.255.255.0

% Subnet not in table
```

To verify the configuration:

On BB1

```
BB1#Show mpls ldp discovery all

Local LDP Identifier:
    7.7.7.7:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 6.6.6.6:0
```

Note the message stating "no host route" is no longer there.

```
BB1#Show mpls ldp discovery detail

Local LDP Identifier:
    7.7.7.7:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 7.7.7.7
            LDP Id: 6.6.6.6:0
              Src IP addr: 10.1.67.6; Transport IP addr: 6.6.6.6
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Reachable via 6.6.6.6/32
```

Note the "no host route to transport addr" is no longer there.

On R1

```
R1#Show mpls ldp discovery all

Local LDP Identifier:
    1.1.1.1:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 2.2.2.2:0
```

On R2

```
R2#Show mpls ldp discovery all

Local LDP Identifier:
    2.2.2.2:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 1.1.1.1:0
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 3.3.3.3:0
```

On R3

```
R3#Show mpls ldp discovery all

Local LDP Identifier:
    3.3.3.3:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 2.2.2.2:0
        Serial0/0.34 (ldp): xmit/recv
            LDP Id: 4.4.4.4:0
        Serial0/0.35 (ldp): xmit/recv
            LDP Id: 5.5.5.5:0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4#Show mpls ldp discovery all

Local LDP Identifier:
    4.4.4.4:0
    Discovery Sources:
    Interfaces:
        Serial0/0.43 (ldp): xmit/recv
            LDP Id: 3.3.3.3:0
        Serial0/0.45 (ldp): xmit/recv
            LDP Id: 5.5.5.5:0
```

On R5

```
R5#Show mpls ldp discovery all

Local LDP Identifier:
    5.5.5.5:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 6.6.6.6:0
        Serial0/0.53 (ldp): xmit/recv
            LDP Id: 3.3.3.3:0
        Serial0/0.54 (ldp): xmit/recv
            LDP Id: 4.4.4.4:0
```

On R6

```
R6#Show mpls ldp discovery all

Local LDP Identifier:
    6.6.6.6:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 7.7.7.7:0
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 5.5.5.5:0
```

On BB1

```
BB1#Show mpls ldp discovery all

Local LDP Identifier:
    7.7.7.7:0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 6.6.6.6:0
```

NOTE: The "NO host route to transport addr" is informational ONLY, in this case it will NOT cause a problem. The LDP router-id and the transport address are two different issues, if the transport address is NOT specifically configured, it will default to the router-id.

The "MPLS ldp discovery transport-address interface" interface configuration mode command can be used to change the transport IP address to be the interface's IP address instead of the loopback0 interface.

The following demonstrates this fact:

On BB1

```
BB1(config)#int F0/0
BB1(config-if)#MPLS ldp discovery transport-address interface
```

On R6

```
R6(config)#int F0/0
R6(config-if)#mpls ldp discovery transport-address interface
```

To verify the configuration:

On BB1

```
BB1#Show mpls ldp discovery det

 Local LDP Identifier:
    7.7.7.7:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 10.1.67.7
            LDP Id: 6.6.6.6:0; no host route to transport addr
              Src IP addr: 10.1.67.6; Transport IP addr: 10.1.67.6
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Reachable via 10.1.67.0/24
```

On R6

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R6#Show mpls ldp discovery det

Local LDP Identifier:
    6.6.6.6:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 10.1.67.6
            LDP Id: 7.7.7.7:0; no host route to transport addr
              Src IP addr: 10.1.67.7; Transport IP addr: 10.1.67.7
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Reachable via 10.1.67.0/24
        FastEthernet0/1 (ldp): xmit/recv
            Enabled: Interface config
            Hello interval: 5000 ms; Transport IP addr: 6.6.6.6
            LDP Id: 5.5.5.5:0
              Src IP addr: 10.1.56.5; Transport IP addr: 5.5.5.5
              Hold time: 15 sec; Proposed local/peer: 15/15 sec
              Reachable via 5.5.5.5/32
```

Task 3

Configure the interval of discovery hellos to be 15 seconds with a hold timer of 45 seconds on all LSRs.

Remember that LSRs will send discovery hellos every 5 seconds with a holdtimer of 15 seconds; these are sent to all routers on local subnets using UDP port 646.

The following shows the different ways to see the default parameters:

On BB1

```
BB1#Show mpls ldp discovery det | Inc Hello|Hold

            Hello interval: 5000 ms; Transport IP addr: 10.1.67.7
              Hold time: 15 sec; Proposed local/peer: 15/15 sec

BB1#Show mpls ldp parameters | Inc Discovery hello

Discovery hello: holdtime: 15 sec; interval: 5 sec
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To change the timers based on the requirements of this task:

On BB1

```
BB1(config)#mpls ldp discovery hello holdtime 45
BB1(config)#mpls ldp discovery hello interval 15
```

To verify the configuration:

On BB1

Note since the timers of R6 have not changed yet, they are set to their default value, but the local hold time is set to 45 seconds.

```
BB1#Show mpls ldp discovery det | Inc Hello|Hold

            Hello interval: 5000 ms; Transport IP addr: 10.1.67.7
              Hold time: 15 sec; Proposed local/peer: 45/15 sec
```

The output of the following show command reveals that the local timers are set to 15/45.

```
BB1#Show mpls ldp parameters | Inc Discovery hello

Discovery hello: holdtime: 45 sec; interval: 15 sec
```

On R6

```
R6(config)#mpls ldp discovery hello interval 15
R6(config)#mpls ldp discovery hello holdtime 45
```

To verify the configuration:

On R6

```
R6#Sh mpls ldp discovery det

            Hello interval: 15000 ms; Transport IP addr: 10.1.67.6
              Hold time: 45 sec; Proposed local/peer: 45/45 sec
            Hello interval: 5000 ms; Transport IP addr: 6.6.6.6
              Hold time: 15 sec; Proposed local/peer: 45/15 sec
```

On BB1

```
BB1#Show mpls ldp discovery det | Inc Hello|Hold
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Hello interval: 15000 ms; Transport IP addr: 10.1.67.7
  Hold time: 45 sec; Proposed local/peer: 45/45 sec
```

On All routers:

```
Rx(config)#mpls ldp discovery hello interval 15
Rx(config)#mpls ldp discovery hello holdtime 45
```

Task 4

Configure the session keepalives and hold timers of all routers to 30 and 90 seconds respectively.

On BB1

The following show command displays the default values of session keepalives and hold timers. When an LDP session is established between two LSRs, the hold time used for the session is the lower of the values configured on the two LSRs.

```
BB1#Show mpls ldp parameters | Inc Session

Session hold time: 180 sec; keep alive interval: 60 sec
```

To change the keepalives to 30 seconds and the hold timer to 90 seconds:

```
BB1(config)#mpls ldp holdtime 90
```

To verify the configuration:

On BB1

```
BB1#Show mpls ldp parameters  | Inc Session

Session hold time: 90 sec; keep alive interval: 30 sec
```

Note once the hold time is changed, the keepalives are set to roughly 1/3 of the hold timer.

On All Routers:

```
Rx(config)#mpls ldp holdtime 90
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Configure the router-id of R1 to be its Loopback 1 interface. You should NOT reload the router to accomplish this task.

Note the LDP-ID can be found using the following show commands:

On R1

```
R1#Show mpls ldp discovery

Local LDP Identifier:
    1.1.1.1:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 2.2.2.2:0
```

On R2

```
R2#Show mpls ldp discovery

Local LDP Identifier:
    2.2.2.2:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 1.1.1.1:0
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 3.3.3.3:0
```

Note, once the LDP-ID is advertised in the IGP, the LDP-ID can be changed, the LDP-ID is changed using the "MPLS LDP router-id" global configuration command, if LDP-ID is changed, a reload must occur before the LSR assumes the newly configured router-id, unless the "force" keyword is used with the command.

On R1

```
R1(config)#mpls ldp router-id loopback 1 force
```

The keyword "force" resets the tcp session and uses the new router-id. You should see the following console messages:

%LDP-5-NBRCHG: LDP Neighbor 2.2.2.2:0 (1) is DOWN (LDP Router ID changed)

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

NOTE: The console message states that the LDP Neighbor 2.2.2.2:0 is down, but when would it come up?

Let's check:

On R1

```
R1#Show mpls ldp Neighbor
R1#
```

R1 does NOT have any neighbors.

```
R1#Show mpls ldp discovery

 Local LDP Identifier:
    11.1.1.1:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 2.2.2.2:0
```

If R1 does not see R2 as a neighbor, then, how does it see 2.2.2.2:0 in the above output? Remember that the above show command ONLY states the router that the local router is exchanging LDP discovery hello messages with; the Hello messages are sent to 224.0.0.2, and as long as R2 is up, it should process the messages. But to establish an LDP session, TCP 646 is used and R2 should have reachability to the neighbor's transport address, in this case the RID of R1. Let's check the routing table of R2:

On R2

```
R2#Show ip route ospf | Inc O

O       1.1.1.1 [110/2] via 10.1.12.1, 01:15:59, FastEthernet0/0
O       3.3.3.3 [110/2] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       4.4.4.4 [110/66] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       5.5.5.5 [110/66] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       6.6.6.6 [110/67] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       7.7.7.7 [110/68] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       10.1.45.0 [110/129] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       10.1.35.0 [110/65] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       10.1.34.0 [110/65] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       10.1.56.0 [110/66] via 10.1.23.3, 01:15:59, FastEthernet0/1
O       10.1.67.0 [110/67] via 10.1.23.3, 01:15:59, FastEthernet0/1
```

You can see that the local router (R2) does NOT have a route to R1's RID. Let's advertise the

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

loopback1 interface of R1 in OSPF:

On R1

```
R1(config)#Router ospf 1
R1(config-router)#Netw 11.1.1.1 0.0.0.0 area 0
```

NOTE: As soon as the Loopback1 interface of R1 was advertised the LDP Neighbor was established, and the following console message was displayed:

%LDP-5-NBRCHG: LDP Neighbor 2.2.2.2:0 (1) is UP

```
R1#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 11.1.1.1:0
        TCP connection: 2.2.2.2.646 - 11.1.1.1.50946
        State: Oper; Msgs sent/rcvd: 23/25; Downstream
        Up time: 00:02:50
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.12.2
        Addresses bound to peer LDP Ident:
          10.1.12.2       10.1.23.2       2.2.2.2         22.2.2.2

R2#Show mpls ldp neighbor | B 11.1.1.1

    Peer LDP Ident: 11.1.1.1:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 11.1.1.1.50946 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 29/27; Downstream
        Up time: 00:04:32
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.12.1
        Addresses bound to peer LDP Ident:
          10.1.12.1       1.1.1.1         11.1.1.1
```

Task 6

Configure authentication between R1 and R2 using "**Cisco**" as the password.

Authentication can be invoked between two LDP peers to verify each segment sent on the TCP connection between the peers. The authentication must be configured on both peers using the same password.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1(config)#mpls ldp neighbor 2.2.2.2 password Cisco
```

Note once R1 is configured, the session between R1 and R2 will be torn down and the following console message will be displayed:

%LDP-5-NBRCHG: LDP Neighbor 2.2.2.2:0 is DOWN (Session's MD5 password changed)
%TCP-6-BADAUTH: No MD5 digest from 2.2.2.2(646) to 11.1.1.1(30627)

On R2

```
R2(config)#mpls ldp neighbor 11.1.1.1 password Cisco
```

You should get the following message:

%LDP-5-NBRCHG: LDP Neighbor 11.1.1.1:0 is UP

Task 7

The label space of routers is platform dependent. By default, the routers begin numbering the labels with 16 up to 100,000. Change the label space such that the routers use the following labels:

Router	Label range:
R1	100 - 199
R2	200 – 299
R3	300 – 399
R4	400 – 499
R5	500 – 599
R6	600 – 699
BB1	700 – 799

Use the following show command to see the default values, note the minimum value in the range is 16 and the maximum value within the range is 100,000. Therefore, the LSR will start assigning labels starting from 16 and it will go up to 100,000, which means that the LSR can assign up to 99984 labels.

On R1

```
R1#Sh mpls ldp parameters | Inc label
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Downstream label generic region: min label: 16; max label: 100000
```

You would not normally change a label range after MPLS is fully deployed. The planning of labels should be performed during the initial phase of implementation. But if they must be changed, a reload is required for the new range to take effect.

The following command changes the label range:

On R1

```
R1(config)#mpls label range 100 199
```

You should receive the following message:

```
% Label range changes will take effect at the next reload.
```

To verify the configuration:

On R1

```
R1#Sh mpls ldp parameters I Inc label

Downstream label generic region: min label: 16; max label: 100000
   [Configured: min label: 100; max label: 199]
```

Note the above Show command states that the label range should be from 100 to 199 after a reload; this is why the configured range is between brackets.

To verify the configuration:

On R1

```
R1#wr
R1#Reload
```

When the router comes up:

```
R1#Sh mpls ldp parameters | Inc label

Downstream label generic region: min label: 100; max label: 199
```

On R2

```
R2(config)#mpls label range 200 299
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#mpls label range 300 399
```

On R4

```
R4(config)#mpls label range 400 499
```

On R5

```
R5(config)#mpls label range 500 599
```

On R6

```
R6(config)#mpls label range 600 699
```

On BB1

```
BB1(config)#mpls label range 700 799
```

On All Routers

```
Rx#Wr
Rx#Reload
```

To see how labels are generated and exchanged and what's used hop by hop:

On R1

```
R1#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.12.2, 00:11:03, FastEthernet0/0
O       3.3.3.3 [110/3] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       4.4.4.4 [110/67] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       5.5.5.5 [110/67] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       6.6.6.6 [110/68] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       7.7.7.7 [110/69] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       10.1.23.0 [110/2] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       10.1.45.0 [110/130] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       10.1.35.0 [110/66] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       10.1.34.0 [110/66] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       10.1.56.0 [110/67] via 10.1.12.2, 00:10:43, FastEthernet0/0
O       10.1.67.0 [110/68] via 10.1.12.2, 00:10:43, FastEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip route | Inc C

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
C       1.1.1.0 is directly connected, Loopback0
C       10.1.12.0 is directly connected, FastEthernet0/0
C       11.1.1.0 is directly connected, Loopback1
```

Note the output of the above show command reveals that R1 has 12 OSPF prefixes in its routing table and three directly connected routes for a total of 15 routes.

```
R1#Show mpls ldp binding local

  tib entry: 1.1.1.0/24, rev 4
        local binding:  tag: imp-null
  tib entry: 2.2.2.2/32, rev 18
        local binding:  tag: 105
  tib entry: 3.3.3.3/32, rev 16
        local binding:  tag: 104
  tib entry: 4.4.4.4/32, rev 14
        local binding:  tag: 103
  tib entry: 5.5.5.5/32, rev 12
        local binding:  tag: 102
  tib entry: 6.6.6.6/32, rev 10
        local binding:  tag: 101
  tib entry: 7.7.7.7/32, rev 8
        local binding:  tag: 100
  tib entry: 10.1.12.0/24, rev 5
        local binding:  tag: imp-null
  tib entry: 10.1.23.0/24, rev 30
        local binding:  tag: 111
  tib entry: 10.1.34.0/24, rev 24
        local binding:  tag: 108
  tib entry: 10.1.35.0/24, rev 22
        local binding:  tag: 107
  tib entry: 10.1.45.0/24, rev 20
        local binding:  tag: 106
  tib entry: 10.1.56.0/24, rev 28
        local binding:  tag: 110
  tib entry: 10.1.67.0/24, rev 26
        local binding:  tag: 109

  tib entry: 11.1.1.0/24, rev 6
        local binding:  tag: imp-null
```

Note the local router imposed 15 labels, one per prefix; therefore, the routers generate/imposes a label for every entry in their routing table.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The following show command displays the label that R1 generated for prefix 6.6.6.6 /32:

```
R1#Show mpls ldp binding 6.6.6.6 32 local

  tib entry: 6.6.6.6/32, rev 10
        local binding:  tag: 101
```

The following show command diplays the LIB for prefix 6.6.6.6 /32 on R1:

```
R1#Show mpls ldp binding 6.6.6.6 32

  tib entry: 6.6.6.6/32, rev 10
        local binding:  tag: 101
        remote binding: tsr: 2.2.2.2:0, tag: 203
```

Note the output of the above command reveals that R1 imposed label 101 to prefix 6.6.6.6 /32 and it received label 203 from R2 (2.2.2.2)

The following displays the LFIB of R1 for prefix 6.6.6.6 /32:

```
R1#Show mpls forwarding-table 6.6.6.6 32

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
101    203         6.6.6.6/32       0          Fa0/0       10.1.12.2
```

Note The LFIB is what R1 uses to forward labeled traffic toward prefix 6.6.6.6 /32, in this case, R1 generated/imposed label 101 and it swaps this label with label 203 and it will send the traffic out of its F0/0 interface to next hop IP address of 10.1.12.2.

```
R1#Sh ip cef 6.6.6.6

6.6.6.6/32, version 31, epoch 0, cached adjacency 10.1.12.2
0 packets, 0 bytes
  tag information set
    local tag: 101
    fast tag rewrite with Fa0/0, 10.1.12.2, tags imposed: {203}
  via 10.1.12.2, FastEthernet0/0, 0 dependencies
    next hop 10.1.12.2, FastEthernet0/0
    valid cached adjacency
    tag rewrite with Fa0/0, 10.1.12.2, tags imposed: {203}
```

The FIB Table entry also contains the local label information and will impose label "203" on the way out of the Fa0/0 interface. If it should ever receive an unlabeled packet from a peer that is further from the destination network. The LFIB and FIB information MUST compliment each other and contain label info or a packet will not be forwarded correctly or not forwarded at all.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2#Show mpls forwarding 6.6.6.6 32

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
203    300         6.6.6.6/32       0          Fa0/1       10.1.23.3
```

If the local LSR receives a labeled packet with the top LABEL of 203, it will swap this label with 300 and it'll send it out of its F0/1 interface.

On R3

```
R3#Sh mpls forwarding-table 6.6.6.6 32

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
300    505         6.6.6.6/32       0          Se0/0.35    point2point
```

If the local LSR receives a labeled packet with the top LABEL of 300, it will swap this label with 505 and it'll send it out of its S0/0.35 interface.

Let's check the entire LFIB on R4:

On R4

```
R4#Show mpls forwarding-table

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
400    Untagged    5.5.5.5/32       0          Se0/0.45    point2point -Line 1
401    Pop tag     10.1.56.0/24     0          Se0/0.45    point2point -Line 2
402    Pop tag     10.1.35.0/24     0          Se0/0.45    point2point -Line 3
       Pop tag     10.1.35.0/24     0          Se0/0.43    point2point -Line 4
403    504         7.7.7.7/32       0          Se0/0.45    point2point -Line 5
404    505         6.6.6.6/32       0          Se0/0.45    point2point -Line 6
405    Untagged    3.3.3.3/32       0          Se0/0.43    point2point -Line 7
406    Pop tag     10.1.23.0/24     0          Se0/0.43    point2point -Line 8
407    506         10.1.67.0/24     0          Se0/0.45    point2point -Line 9
408    306         11.1.1.1/32      0          Se0/0.43    point2point -Line 10
409    307         2.2.2.2/32       0          Se0/0.43    point2point -Line 11
410    308         1.1.1.1/32       0          Se0/0.43    point2point -Line 12
411    309         10.1.12.0/24     0          Se0/0.43    point2point -Line 13
```

Explanation of the above output:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Line 1: (Line 7 is identical)
If the local LSR receives a labeled packet with a top LABEL of 400, it will remove all labels and forward the packet as an IP packet, because the outgoing Tag is "Untagged", and it will forward it on S0/0.45 interface. This is label to IP forwading.

Line 2: (Line 8 is identical)
If the local LSR receives a labeled packet with a top LABEL of 401, it will remove the top label (Pop one label) "Pop tag" and forwards the packet as labeled or as an IP packet on S0/0.45 interface. This could be label to IP forwarding or label to label forwarding.

Line 3 and 4:
If the local LSR receives a labeled packet with a top LABEL of 402, it would remove the top label (Pop one label) "Pop tag" and forwards the packet as labeled or as an IP packet on S0/0.45, but in this case notice how the enrty below it has no label, but the entry is for the same prefix, this means that load balancing is being performed, and if the routing/CEF table is checked for this prefix, you should see two entries for this prefix, as follows:

```
R4#Sh ip cef 10.1.35.0

10.1.35.0/24, version 33, epoch 0, per-destination sharing
0 packets, 0 bytes
  tag information set
    local tag: 402
  via 10.1.45.5, Serial0/0.45, 0 dependencies
    traffic share 1
    next hop 10.1.45.5, Serial0/0.45
    valid adjacency
    tag rewrite with Se0/0.45, point2point, tags imposed: {}
  via 10.1.34.3, Serial0/0.43, 0 dependencies
    traffic share 1
    next hop 10.1.34.3, Serial0/0.43
    valid adjacency
    tag rewrite with Se0/0.43, point2point, tags imposed: {}
  0 packets, 0 bytes switched through the prefix
  tmstats: external 0 packets, 0 bytes
           internal 0 packets, 0 bytes


R4#Show ip route 10.1.35.0 255.255.255.0

Routing entry for 10.1.35.0/24
  Known via "ospf 1", distance 110, metric 128, type intra area
  Last update from 10.1.34.3 on Serial0/0.43, 01:30:21 ago
  Routing Descriptor Blocks:
  * 10.1.45.5, from 55.5.5.5, 01:30:31 ago, via Serial0/0.45
      Route metric is 128, traffic share count is 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    10.1.34.3, from 33.3.3.3, 01:30:21 ago, via Serial0/0.43
      Route metric is 128, traffic share count is 1
```

Line 5: (Lines 6, and 9-13 are identical)

If the local LSR receives a labeled packet with the top LABEL of 403, it will swap the label with 504 and then it will forward it on the interface specified in the outgoing interface column.

Therefore, when the local router receives a labeled packet with the top label of 404, it will swap the label with 505 and forward the packet out of its S0/0.45. Let's check R5:

On R5

```
R5#Show mpls forwarding-table 6.6.6.6 32

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
505    Untagged    6.6.6.6/32       0          Fa0/1       10.1.56.6
```

When the local router (R5) receives a labeled packet with a top label of 505, it will remove all labels and forward the packet out of its F0/1 interface.

On R6

```
R6#Show mpls forwarding-table 6.6.6.6 32

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface

R6#Show ip route 6.6.6.6

Routing entry for 6.6.6.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:

  * directly connected, via Loopback0
      Route metric is 0, traffic share count is 1
```

Task 8

Configure F0/1 interface of BB1 with an IP address of 177.7.7.7 /24 and advertise this interface in OSPF area 0.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

On BB1

```
BB1(config)#int F0/1
BB1(config-if)#ip addr 177.7.7.7 255.255.255.0
BB1(config-if)#NO Keep
BB1(config-if)#NO shut

BB1(config-if)#Router ospf 1
BB1(config-router)#Netw 177.7.7.7 0.0.0.0 area 0
```

Task 9

Starting from BB1 examine the Control Plane for network 177.7.7.0 /24 UP/stream toward R1.

On BB1

```
BB1#Show mpls interfaces

Interface              IP             Tunnel   Operational
FastEthernet0/0        Yes (ldp)      No       Yes
```

The F0/0 interface on BB1 has been configured for mpls ldp and its operational. The Tunnel section should state "NO" due to the fact that Traffic Engineering has NOT been configured for this tunnel.

```
BB1#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 40
        local binding:  tag: imp-null
        remote binding: tsr: 6.6.6.6:0, tag: 612
```

The local binding is "implicit null". This is the default behavior for all directly connected routes in the routing table. R6 is considered the Up/stream peer for the tunnel. This means that the route resides here locally.

The remote binding will not be used. OSPF will not choose to route to a local route through R6.

On R6

```
R6#Sh mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 40
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        local binding:  tag: 612
        remote binding: tsr: 7.7.7.7:0, tag: imp-null
        remote binding: tsr: 5.5.5.5:0, tag: 511
```

Note R6 received two remote bindings 511 from neighbor 5.5.5.5 and "Implicit-Null" from BB1.

On R5

```
R5#Sh mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 41
        local binding:  tag: 511
        remote binding: tsr: 3.3.3.3:0, tag: 311
        remote binding: tsr: 6.6.6.6:0, tag: 612
        remote binding: tsr: 4.4.4.4:0, tag: 412
```



Note in this case R5 received three remote binding for this prefix, but which one will it use?
Well….in this case the LFIB of R5 should be consulted:

On R5

```
R5#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
511    612         177.7.7.0/24     0          Fa0/1       10.1.56.6
```



On R4

```
R4#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 39

        local binding:  tag: 412
        remote binding: tsr: 3.3.3.3:0, tag: 311
        remote binding: tsr: 5.5.5.5:0, tag: 511
```





```
R4#Show mpls forwarding-table 177.7.7.0 24

Local   Outgoing     Prefix           Bytes tag  Outgoing    Next Hop

tag     tag or VC   or Tunnel Id     switched   interface
412     511         177.7.7.0/24     0          Se0/0.45    point2point
```





On R3

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
R3#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 41
        local binding:  tag: 311
        remote binding: tsr: 2.2.2.2:0, tag: 212
        remote binding: tsr: 4.4.4.4:0, tag: 412
        remote binding: tsr: 5.5.5.5:0, tag: 511

R3#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
311    511         177.7.7.0/24     0          Se0/0.35    point2point
```

On R2

```
R2#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 39
        local binding:  tag: 212
        remote binding: tsr: 3.3.3.3:0, tag: 311
        remote binding: tsr: 11.1.1.1:0, tag: 112

R2#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
212    311         177.7.7.0/24     0          Fa0/1       10.1.23.3
```

On R1

```
R1#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 40
        local binding:  tag: 112
        remote binding: tsr: 2.2.2.2:0, tag: 212

R1#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
112    212         177.7.7.0/24     0          Fa0/0       10.1.12.2
```

Note R1 represents the first hop router or the furthest Up/stream router from 177.7.7.0 /24, the local label assigned is 112, this ends the control plane, in the next task the data plane is examined.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 10

Examine the Data Plane for network 177.7.7.0 /24 starting from R1 toward BB1.

On R1

```
R1#Show ip route 177.7.7.0 255.255.255.0
```

Displaying the RIB

```
Routing entry for 177.7.7.0/24

  Known via "ospf 1", distance 110, metric 69, type intra area
  Last update from 10.1.12.2 on FastEthernet0/0, 00:44:54 ago
  Routing Descriptor Blocks:
  * 10.1.12.2, from 77.7.7.7, 00:44:54 ago, via FastEthernet0/0
      Route metric is 69, traffic share count is 1
```

Note the routing table reveals the following:
The route was learned via OSPF process 1
The next hop IP address to reach the destination is 10.1.12.2
The route was advertised by the router with a router id of 77.7.7.7

```
R1#Show ip cef 177.7.7.0 255.255.255.0
```

Displaying the FIB

```
177.7.7.0/24, version 41, epoch 0, cached adjacency 10.1.12.2
0 packets, 0 bytes
  tag information set
    local tag: 112
    fast tag rewrite with Fa0/0, 10.1.12.2, tags imposed: {212}
  via 10.1.12.2, FastEthernet0/0, 0 dependencies
    next hop 10.1.12.2, FastEthernet0/0
    valid cached adjacency
    tag rewrite with Fa0/0, 10.1.12.2, tags imposed: {212}
```

Note the FIB gets built based on the RIB, in this case if an unlabeled packet is received by the local router destined for network 177.7.7.0 /24 the router will use FIB to forward the route, and a label of 212 will be imposed before the packet leaves the router, this router is the "Ingress Edge LSR".

```
R1#Show mpls ldp binding 177.7.7.0 24
```

Displaying the LIB

```
  tib entry: 177.7.7.0/24, rev 40
        local binding:  tag: 112
        remote binding: tsr: 2.2.2.2:0, tag: 212
```

```
R1#Show mpls forwarding-table 177.7.7.0 24
```

Displaying the LFIB

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
112    212         177.7.7.0/24     0          Fa0/0       10.1.12.2
```

If the local LSR receives a labeled packet with the top LABEL of 112, it will swap the label with 212, and then it will forward it out of the interface specified in the outgoing interface column (F0/0 in this case).

On R2

```
R2#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 39
        local binding:  tag: 212
        remote binding: tsr: 3.3.3.3:0, tag: 311
        remote binding: tsr: 11.1.1.1:0, tag: 112

R2#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
212    311         177.7.7.0/24     0          Fa0/1       10.1.23.3
```

If the local LSR receives a labeled packet with the top LABEL of 212 it will swap the label with 311, and then it will forward it out of the interface specified in the outgoing interface column (In this case F0/1).

On R3

```
R3#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 41
        local binding:  tag: 311
        remote binding: tsr: 2.2.2.2:0, tag: 212
        remote binding: tsr: 4.4.4.4:0, tag: 412
        remote binding: tsr: 5.5.5.5:0, tag: 511


R3#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
311    511         177.7.7.0/24     0          Se0/0.35    point2point
```

If the local LSR receives a labeled packet with the top LABEL of 311 it will swap the label with 511,

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

and then it will forward it out of the interface specified in the outgoing interface column (In this case S0/0.35).

```
R3#Show mpls ip binding 177.7.7.0 24

  177.7.7.0/24
        in label:      311
        out label:     212        lsr: 2.2.2.2:0
        out label:     412        lsr: 4.4.4.4:0
        out label:     511        lsr: 5.5.5.5:0         inuse
```

Note the above show command is another useful one, because it clearly states that if the local LSR receives a labeled packet with a top label of 311 (in label: 311) it will be swapped with label 511 on the way out, note there are three labels that the local LSR received, but 511 is what's being used (out label: 511 lsr:5.5.5.5:0 inuse).

A very useful command for troubleshooting.

On R4

```
R4#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 39
        local binding:  tag: 412
        remote binding: tsr: 3.3.3.3:0, tag: 311
        remote binding: tsr: 5.5.5.5:0, tag: 511

R4#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
412    511         177.7.7.0/24     0          Se0/0.45    point2point
```

If the local LSR receives a labeled packet with the top LABEL of 412 it will swap the label with 511, and then it will forward it out of the interface specified in the outgoing interface column (In this case S0/0.45).

```
R4#Show mpls ip binding 177.7.7.0 24

  177.7.7.0/24
        in label:      412
        out label:     311        lsr: 3.3.3.3:0
        out label:     511        lsr: 5.5.5.5:0         inuse
```

On R5

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 41
        local binding:  tag: 511
        remote binding: tsr: 3.3.3.3:0, tag: 311
        remote binding: tsr: 6.6.6.6:0, tag: 612
        remote binding: tsr: 4.4.4.4:0, tag: 412

R5#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
511    612         177.7.7.0/24     0          Fa0/1       10.1.56.6
```

If the local LSR receives a labeled packet with the top LABEL of 511 it will swap the label with 612, and then it will forward it out of the interface specified in the outgoing interface column (In this case F0/1).

```
R5#Show mpls ip binding 177.7.7.0 24

  177.7.7.0/24
        in label:     511
        out label:    311        lsr: 3.3.3.3:0
        out label:    612        lsr: 6.6.6.6:0        inuse
        out label:    412        lsr: 4.4.4.4:0
```

On R6

```
R6#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 40
        local binding:  tag: 612
        remote binding: tsr: 7.7.7.7:0, tag: imp-null
        remote binding: tsr: 5.5.5.5:0, tag: 511

R6#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
612    Pop tag     177.7.7.0/24     0          Fa0/0       10.1.67.7
```

If the local LSR receives a labeled packet with a top LABEL of 612, it will remove the top label (Pop one label) "Pop tag" and forwards the packet as labeled or as an IP packet out of F0/0 interface. The implicit-null can be thought of as a set of instructions that tells the seond last router to pop/remove the label; label 3 is used for this purpose.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Remember that the local LSR is one hop prior to the last (Destination) and it's called the "Penultimate Hop Popper". By performing PHP (Penultimate Hop Popping), the local LSR will NOT send a labeled packet to BB1 and as a result of that BB1 won't have to perform 2 table lookups (LFIB to remove the label and FIB to know where to send the packet), and instead, BB1 will receive an unlabeled packet, in which case it will forward the packet based on the routing table.

But what if the packet contains two labels and the second label is a QOS label?
Well….in this case the second last router (PHP) should NOT remove all the labels. This is when the "Explicit-Null label" or label number 0 is used.

```
R6#Show mpls ip binding 177.7.7.0 24

  177.7.7.0/24
        in label:     612
        out label:    imp-null  lsr: 7.7.7.7:0        inuse
        out label:    511       lsr: 5.5.5.5:0
```

On BB1

```
BB1#Show mpls ldp binding 177.7.7.0 24

  tib entry: 177.7.7.0/24, rev 40
        local binding:  tag: imp-null
        remote binding: tsr: 6.6.6.6:0, tag: 612

BB1#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface

BB1#Show ip route 177.7.7.0 255.255.255.0

Routing entry for 177.7.7.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via FastEthernet0/1
      Route metric is 0, traffic share count is 1
```

Note when the local LSR receives a packet destined for 177.7.7.0 /24 which is directly connected, it will forward it out of F0/1 interface. Note the packet is unlabeled, therefore, the local LSR won't have to perform two lookups, and it will simply check the routing table and forward it out of F0/1 interface.

BB1 can be configured to advertise an "Explicit-Null" to R6 as follows:

On BB1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
BB1(config)#Int F0/0
BB1(config-if)#NO mpls ip
BB1(config-if)#Exit

BB1(config)#Mpls LDP explicit-null

BB1(config)#Int F0/0
BB1(config-if)#mpls ip
```

You should see the following console message stating that you have establishd an LDP session with an LSR with a RID of 6.6.6.6:

```
%LDP-5-NBRCHG: LDP Neighbor 6.6.6.6:0 (1) is UP
```

To verify the configuration:

On R6

```
R6#Show mpls forwarding-table 177.7.7.0 24

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
612    0           177.7.7.0/24     0          Fa0/0       10.1.67.7

R6#Show mpls ip binding 177.7.7.0 24

  177.7.7.0/24
        in label:      612
        out label:     511         lsr: 5.5.5.5:0
        out label:     exp-null    lsr: 7.7.7.7:0        inuse
```

Normally, LDP advertises an Implicit null label for all directly connected routes, which instructs the next hop router, to perform PHP. But if you require the next hop router to replace the incoming label with an "Explicit null" label, then, the "Mpls ldp explicit-null" global configuration command can be used.

This command has two keywords a "for" and a "to".

The "for" keyword: Specifies prefixes FOR which the explicit-null should be advertised.

The "to" keyword: Specifies LDP peers to which Explicit Null should be advertised TO.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 11

Configure LDP Conditional Label Advertising to exclude the links that interconnect the routers in this topology.

Step 1:
Configure R1 to stop advertising labels all together:

On R1

```
R1(config)#NO mpls ldp advertise-labels
```

To verify the configuration:

On R2

```
R2#Sh mpls ldp binding | S remote

        remote binding: tsr: 3.3.3.3:0, tag: 308
  tib entry: 2.2.2.0/24, rev 5
        remote binding: tsr: 3.3.3.3:0, tag: 307
  tib entry: 3.3.3.0/24, rev 37
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 3.3.3.3/32, rev 30
        remote binding: tsr: 3.3.3.3:0, tag: 302
  tib entry: 5.5.5.5/32, rev 28
        remote binding: tsr: 3.3.3.3:0, tag: 301
  tib entry: 6.6.6.6/32, rev 27
        remote binding: tsr: 3.3.3.3:0, tag: 300
  tib entry: 7.7.7.7/32, rev 26
        remote binding: tsr: 3.3.3.3:0, tag: 310
  tib entry: 10.1.12.0/24, rev 7
        remote binding: tsr: 3.3.3.3:0, tag: 309
  tib entry: 10.1.23.0/24, rev 8
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 10.1.34.0/24, rev 33
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 10.1.35.0/24, rev 32
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 10.1.45.0/24, rev 31
        remote binding: tsr: 3.3.3.3:0, tag: 304
  tib entry: 10.1.56.0/24, rev 35
        remote binding: tsr: 3.3.3.3:0, tag: 305
  tib entry: 10.1.67.0/24, rev 34
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        remote binding: tsr: 3.3.3.3:0, tag: 303
  tib entry: 11.1.1.1/32, rev 10
        remote binding: tsr: 3.3.3.3:0, tag: 306
  tib entry: 22.2.2.0/24, rev 6
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 39
        remote binding: tsr: 3.3.3.3:0, tag: 311
```

Note the output of the above show command reveals that R2 is no longer receiving any labels from R1.

Step 2:

In this step an access-list is configured to deny the networks assigned to the links that interconnect the LSRs and allow everything else, this access-list will be referenced by the "FOR" keyword of the "mpls ldp advertise-labels" command:

On R1

```
R1(config)#access-list 1 deny 10.1.12.0 0.0.0.255
R1(config)#access-list 1 deny 10.1.23.0 0.0.0.255
R1(config)#access-list 1 deny 10.1.34.0 0.0.0.255
R1(config)#access-list 1 deny 10.1.35.0 0.0.0.255
R1(config)#access-list 1 deny 10.1.45.0 0.0.0.255
R1(config)#access-list 1 deny 10.1.56.0 0.0.0.255
R1(config)#access-list 1 deny 10.1.67.0 0.0.0.255
R1(config)#access-list 1 permit any
```

Step 3:

In this step an access-list is configured to identify the peer/s that the label/s will be advertised to; this access-list will be referenced by the "TO" keyword of the "mpls ldp advertise-labels" command:

```
R1(config)#Access-list 2 permit any
```

Step 4:

The "mpls ldp advertise-labels" command is configured referencing the two access-lists:

```
R1(config)#mpls ldp advertise-labels for 1 to 2
```

Step 5:

Verify:

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

On R2

```
R2#Sh mpls ldp binding neighbor 1.1.1.1

  tib entry: 1.1.1.0/24, rev 40
        remote binding: tsr: 11.1.1.1:0, tag: imp-null
  tib entry: 2.2.2.2/32, rev 20
        remote binding: tsr: 11.1.1.1:0, tag: 105
  tib entry: 3.3.3.3/32, rev 30
        remote binding: tsr: 11.1.1.1:0, tag: 104
  tib entry: 4.4.4.4/32, rev 29
        remote binding: tsr: 11.1.1.1:0, tag: 103
  tib entry: 5.5.5.5/32, rev 28
        remote binding: tsr: 11.1.1.1:0, tag: 102
  tib entry: 6.6.6.6/32, rev 27
        remote binding: tsr: 11.1.1.1:0, tag: 101
  tib entry: 7.7.7.7/32, rev 26
        remote binding: tsr: 11.1.1.1:0, tag: 100
  tib entry: 11.1.1.0/24, rev 41
        remote binding: tsr: 11.1.1.1:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 39
        remote binding: tsr: 11.1.1.1:0, tag: 112
```

Note R1 (1.1.1.1) is no longer advertising the links to this neighbor.

On ALL Routers:

```
Rx(config)#NO mpls ldp advertise-labels

Rx(config)#access-list 1 deny 10.1.12.0 0.0.0.255
Rx(config)#access-list 1 deny 10.1.23.0 0.0.0.255
Rx(config)#access-list 1 deny 10.1.34.0 0.0.0.255
Rx(config)#access-list 1 deny 10.1.35.0 0.0.0.255
Rx(config)#access-list 1 deny 10.1.45.0 0.0.0.255
Rx(config)#access-list 1 deny 10.1.56.0 0.0.0.255
Rx(config)#access-list 1 deny 10.1.67.0 0.0.0.255
Rx(config)#access-list 1 permit any

Rx(config)#Access-list 2 permit any

Rx(config)#mpls ldp advertise-labels for 1 to 2
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On BB1

```
BB1#Show mpls ldp binding neighbor 6.6.6.6

  tib entry: 1.1.1.1/32, rev 58
        remote binding: tsr: 6.6.6.6:0, tag: 610
  tib entry: 2.2.2.2/32, rev 59
        remote binding: tsr: 6.6.6.6:0, tag: 609
  tib entry: 3.3.3.3/32, rev 60
        remote binding: tsr: 6.6.6.6:0, tag: 602
  tib entry: 4.4.4.4/32, rev 61
        remote binding: tsr: 6.6.6.6:0, tag: 601
  tib entry: 5.5.5.5/32, rev 62
        remote binding: tsr: 6.6.6.6:0, tag: 600
  tib entry: 6.6.6.0/24, rev 68
        remote binding: tsr: 6.6.6.6:0, tag: imp-null
  tib entry: 7.7.7.7/32, rev 69
        remote binding: tsr: 6.6.6.6:0, tag: 607
  tib entry: 11.1.1.1/32, rev 65
        remote binding: tsr: 6.6.6.6:0, tag: 608
  tib entry: 66.6.6.0/24, rev 70
        remote binding: tsr: 6.6.6.6:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 67
        remote binding: tsr: 6.6.6.6:0, tag: 612
```

On R6

```
R6#Show mpls ldp binding neighbor 5.5.5.5

  tib entry: 1.1.1.1/32, rev 65
        remote binding: tsr: 5.5.5.5:0, tag: 509
  tib entry: 2.2.2.2/32, rev 66
        remote binding: tsr: 5.5.5.5:0, tag: 508
  tib entry: 3.3.3.3/32, rev 67
        remote binding: tsr: 5.5.5.5:0, tag: 501
  tib entry: 4.4.4.4/32, rev 68
        remote binding: tsr: 5.5.5.5:0, tag: 500
  tib entry: 5.5.5.0/24, rev 75
        remote binding: tsr: 5.5.5.5:0, tag: imp-null
  tib entry: 6.6.6.6/32, rev 76
        remote binding: tsr: 5.5.5.5:0, tag: 505
  tib entry: 7.7.7.7/32, rev 71
        remote binding: tsr: 5.5.5.5:0, tag: 504
  tib entry: 11.1.1.1/32, rev 72
        remote binding: tsr: 5.5.5.5:0, tag: 507
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  tib entry: 55.5.5.0/24, rev 77
        remote binding: tsr: 5.5.5.5:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 74
        remote binding: tsr: 5.5.5.5:0, tag: 511

R6#Show mpls ldp binding neighbor 77.7.7.7

  tib entry: 1.1.1.1/32, rev 65
        remote binding: tsr: 7.7.7.7:0, tag: 713
  tib entry: 2.2.2.2/32, rev 66
        remote binding: tsr: 7.7.7.7:0, tag: 714
  tib entry: 3.3.3.3/32, rev 67
        remote binding: tsr: 7.7.7.7:0, tag: 715
  tib entry: 4.4.4.4/32, rev 68
        remote binding: tsr: 7.7.7.7:0, tag: 716
  tib entry: 5.5.5.5/32, rev 69
        remote binding: tsr: 7.7.7.7:0, tag: 717
  tib entry: 6.6.6.6/32, rev 76
        remote binding: tsr: 7.7.7.7:0, tag: 718
  tib entry: 7.7.7.0/24, rev 78
        remote binding: tsr: 7.7.7.7:0, tag: exp-null
  tib entry: 11.1.1.1/32, rev 72
        remote binding: tsr: 7.7.7.7:0, tag: 725
  tib entry: 77.7.7.0/24, rev 79
        remote binding: tsr: 7.7.7.7:0, tag: exp-null
  tib entry: 177.7.7.0/24, rev 74
        remote binding: tsr: 7.7.7.7:0, tag: exp-null
```

On R5

```
R5#Show mpls ldp binding neighbor 6.6.6.6

  tib entry: 1.1.1.1/32, rev 66
        remote binding: tsr: 6.6.6.6:0, tag: 610
  tib entry: 2.2.2.2/32, rev 67
        remote binding: tsr: 6.6.6.6:0, tag: 609
  tib entry: 3.3.3.3/32, rev 68
        remote binding: tsr: 6.6.6.6:0, tag: 602
  tib entry: 4.4.4.4/32, rev 69
        remote binding: tsr: 6.6.6.6:0, tag: 601
  tib entry: 5.5.5.5/32, rev 76
        remote binding: tsr: 6.6.6.6:0, tag: 600
  tib entry: 6.6.6.0/24, rev 77
        remote binding: tsr: 6.6.6.6:0, tag: imp-null
  tib entry: 7.7.7.7/32, rev 72
        remote binding: tsr: 6.6.6.6:0, tag: 607
  tib entry: 11.1.1.1/32, rev 73
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        remote binding: tsr: 6.6.6.6:0, tag: 608
  tib entry: 66.6.6.0/24, rev 78
        remote binding: tsr: 6.6.6.6:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 75
        remote binding: tsr: 6.6.6.6:0, tag: 612

R5#Show mpls ldp binding neighbor 4.4.4.4

  tib entry: 1.1.1.1/32, rev 66
        remote binding: tsr: 4.4.4.4:0, tag: 410
  tib entry: 2.2.2.2/32, rev 67
        remote binding: tsr: 4.4.4.4:0, tag: 409
  tib entry: 3.3.3.3/32, rev 68
        remote binding: tsr: 4.4.4.4:0, tag: 405
  tib entry: 4.4.4.0/24, rev 79
        remote binding: tsr: 4.4.4.4:0, tag: imp-null
  tib entry: 5.5.5.5/32, rev 76
        remote binding: tsr: 4.4.4.4:0, tag: 400
  tib entry: 6.6.6.6/32, rev 71
        remote binding: tsr: 4.4.4.4:0, tag: 404
  tib entry: 7.7.7.7/32, rev 72
        remote binding: tsr: 4.4.4.4:0, tag: 403
  tib entry: 11.1.1.1/32, rev 73
        remote binding: tsr: 4.4.4.4:0, tag: 408
  tib entry: 44.4.4.0/24, rev 80
        remote binding: tsr: 4.4.4.4:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 75
        remote binding: tsr: 4.4.4.4:0, tag: 412

R5#Show mpls ldp binding neighbor 3.3.3.3

  tib entry: 1.1.1.1/32, rev 66
        remote binding: tsr: 3.3.3.3:0, tag: 308
  tib entry: 2.2.2.2/32, rev 67
        remote binding: tsr: 3.3.3.3:0, tag: 307
  tib entry: 3.3.3.0/24, rev 81
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 4.4.4.4/32, rev 69
        remote binding: tsr: 3.3.3.3:0, tag: 302
  tib entry: 5.5.5.5/32, rev 76
        remote binding: tsr: 3.3.3.3:0, tag: 301
  tib entry: 6.6.6.6/32, rev 71
        remote binding: tsr: 3.3.3.3:0, tag: 300
  tib entry: 7.7.7.7/32, rev 72
        remote binding: tsr: 3.3.3.3:0, tag: 310
  tib entry: 11.1.1.1/32, rev 73
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        remote binding: tsr: 3.3.3.3:0, tag: 306
  tib entry: 33.3.3.0/24, rev 82
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 75
        remote binding: tsr: 3.3.3.3:0, tag: 311
```

On R4

```
R4#Show mpls ldp binding neighbor 3.3.3.3

  tib entry: 1.1.1.1/32, rev 62
        remote binding: tsr: 3.3.3.3:0, tag: 308
  tib entry: 2.2.2.2/32, rev 63
        remote binding: tsr: 3.3.3.3:0, tag: 307
  tib entry: 3.3.3.0/24, rev 75
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 4.4.4.4/32, rev 72
        remote binding: tsr: 3.3.3.3:0, tag: 302
  tib entry: 5.5.5.5/32, rev 66
        remote binding: tsr: 3.3.3.3:0, tag: 301
  tib entry: 6.6.6.6/32, rev 67
        remote binding: tsr: 3.3.3.3:0, tag: 300
  tib entry: 7.7.7.7/32, rev 68
        remote binding: tsr: 3.3.3.3:0, tag: 310
  tib entry: 11.1.1.1/32, rev 69
        remote binding: tsr: 3.3.3.3:0, tag: 306
  tib entry: 33.3.3.0/24, rev 76
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 71
        remote binding: tsr: 3.3.3.3:0, tag: 311

R4#Show mpls ldp binding neighbor 5.5.5.5

  tib entry: 1.1.1.1/32, rev 62
        remote binding: tsr: 5.5.5.5:0, tag: 509
  tib entry: 2.2.2.2/32, rev 63
        remote binding: tsr: 5.5.5.5:0, tag: 508
  tib entry: 3.3.3.3/32, rev 64
        remote binding: tsr: 5.5.5.5:0, tag: 501
  tib entry: 4.4.4.4/32, rev 72
        remote binding: tsr: 5.5.5.5:0, tag: 500
  tib entry: 5.5.5.0/24, rev 73
        remote binding: tsr: 5.5.5.5:0, tag: imp-null
  tib entry: 6.6.6.6/32, rev 67
        remote binding: tsr: 5.5.5.5:0, tag: 505
  tib entry: 7.7.7.7/32, rev 68
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        remote binding: tsr: 5.5.5.5:0, tag: 504
  tib entry: 11.1.1.1/32, rev 69
        remote binding: tsr: 5.5.5.5:0, tag: 507
  tib entry: 55.5.5.0/24, rev 74
        remote binding: tsr: 5.5.5.5:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 71
        remote binding: tsr: 5.5.5.5:0, tag: 511
```

On R3

```
R3#Show mpls ldp binding neighbor 4.4.4.4

  tib entry: 1.1.1.1/32, rev 66
        remote binding: tsr: 4.4.4.4:0, tag: 410
  tib entry: 2.2.2.2/32, rev 67
        remote binding: tsr: 4.4.4.4:0, tag: 409
  tib entry: 3.3.3.3/32, rev 76
        remote binding: tsr: 4.4.4.4:0, tag: 405
  tib entry: 4.4.4.0/24, rev 79
        remote binding: tsr: 4.4.4.4:0, tag: imp-null
  tib entry: 5.5.5.5/32, rev 70
        remote binding: tsr: 4.4.4.4:0, tag: 400
  tib entry: 6.6.6.6/32, rev 71
        remote binding: tsr: 4.4.4.4:0, tag: 404
  tib entry: 7.7.7.7/32, rev 72
        remote binding: tsr: 4.4.4.4:0, tag: 403
  tib entry: 11.1.1.1/32, rev 73
        remote binding: tsr: 4.4.4.4:0, tag: 408
  tib entry: 44.4.4.0/24, rev 80
        remote binding: tsr: 4.4.4.4:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 75
        remote binding: tsr: 4.4.4.4:0, tag: 412

R3#Show mpls ldp binding neighbor 5.5.5.5

  tib entry: 1.1.1.1/32, rev 66
        remote binding: tsr: 2.2.2.2:0, tag: 201
  tib entry: 2.2.2.0/24, rev 81
        remote binding: tsr: 2.2.2.2:0, tag: imp-null
  tib entry: 3.3.3.3/32, rev 76
        remote binding: tsr: 2.2.2.2:0, tag: 206
  tib entry: 4.4.4.4/32, rev 69
        remote binding: tsr: 2.2.2.2:0, tag: 205
  tib entry: 5.5.5.5/32, rev 70
        remote binding: tsr: 2.2.2.2:0, tag: 204
  tib entry: 6.6.6.6/32, rev 71
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        remote binding: tsr: 2.2.2.2:0, tag: 203
  tib entry: 7.7.7.7/32, rev 72
        remote binding: tsr: 2.2.2.2:0, tag: 202
  tib entry: 11.1.1.1/32, rev 73
        remote binding: tsr: 2.2.2.2:0, tag: 200
  tib entry: 22.2.2.0/24, rev 82
        remote binding: tsr: 2.2.2.2:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 75
        remote binding: tsr: 2.2.2.2:0, tag: 212
```

On R2

```
R2#Show mpls ldp binding neighbor 3.3.3.3

  tib entry: 1.1.1.1/32, rev 65
        remote binding: tsr: 3.3.3.3:0, tag: 308
  tib entry: 2.2.2.2/32, rev 67
        remote binding: tsr: 3.3.3.3:0, tag: 307
  tib entry: 3.3.3.0/24, rev 77
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 4.4.4.4/32, rev 69
        remote binding: tsr: 3.3.3.3:0, tag: 302
  tib entry: 5.5.5.5/32, rev 70
        remote binding: tsr: 3.3.3.3:0, tag: 301
  tib entry: 6.6.6.6/32, rev 71
        remote binding: tsr: 3.3.3.3:0, tag: 300
  tib entry: 7.7.7.7/32, rev 72
        remote binding: tsr: 3.3.3.3:0, tag: 310
  tib entry: 11.1.1.1/32, rev 74
        remote binding: tsr: 3.3.3.3:0, tag: 306
  tib entry: 33.3.3.0/24, rev 78
        remote binding: tsr: 3.3.3.3:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 76
        remote binding: tsr: 3.3.3.3:0, tag: 311

R2#Show mpls ldp binding neighbor 1.1.1.1

  tib entry: 1.1.1.0/24, rev 64
        remote binding: tsr: 11.1.1.1:0, tag: imp-null
  tib entry: 2.2.2.2/32, rev 67
        remote binding: tsr: 11.1.1.1:0, tag: 105
  tib entry: 3.3.3.3/32, rev 68
        remote binding: tsr: 11.1.1.1:0, tag: 104
  tib entry: 4.4.4.4/32, rev 69
        remote binding: tsr: 11.1.1.1:0, tag: 103
  tib entry: 5.5.5.5/32, rev 70
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        remote binding: tsr: 11.1.1.1:0, tag: 102
  tib entry: 6.6.6.6/32, rev 71
        remote binding: tsr: 11.1.1.1:0, tag: 101
  tib entry: 7.7.7.7/32, rev 72
        remote binding: tsr: 11.1.1.1:0, tag: 100
  tib entry: 11.1.1.0/24, rev 73
        remote binding: tsr: 11.1.1.1:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 76
        remote binding: tsr: 11.1.1.1:0, tag: 112
```

On R1

```
R1#Show mpls ldp binding neighbor 2.2.2.2

   tib entry: 1.1.1.1/32, rev 74
        remote binding: tsr: 2.2.2.2:0, tag: 201
  tib entry: 2.2.2.0/24, rev 75
        remote binding: tsr: 2.2.2.2:0, tag: imp-null
  tib entry: 3.3.3.3/32, rev 65
        remote binding: tsr: 2.2.2.2:0, tag: 206
  tib entry: 4.4.4.4/32, rev 66
        remote binding: tsr: 2.2.2.2:0, tag: 205
  tib entry: 5.5.5.5/32, rev 67
        remote binding: tsr: 2.2.2.2:0, tag: 204
  tib entry: 6.6.6.6/32, rev 68
        remote binding: tsr: 2.2.2.2:0, tag: 203
  tib entry: 7.7.7.7/32, rev 69
        remote binding: tsr: 2.2.2.2:0, tag: 202
  tib entry: 11.1.1.1/32, rev 76
        remote binding: tsr: 2.2.2.2:0, tag: 200
  tib entry: 22.2.2.0/24, rev 77
        remote binding: tsr: 2.2.2.2:0, tag: imp-null
  tib entry: 177.7.7.0/24, rev 73
        remote binding: tsr: 2.2.2.2:0, tag: 212
```

Task 12

In this task the effects of TTL propagation will be tested. Remove the “mpls ip” command from the F0/0 interface of BB1 and R6, and the F0/0 interface of R1 and R2. R1 and BB1 will pose as a customer router that does not have mpls enabled. From BB1 test the connection to 1.1.1.1 using Traceroute.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1 and R2

```
Rx(config)#Int F0/0
Rx(config-if)#NO mpls ip
```

On BB1 and R6

```
(config)#Int F0/0
(config-if)#NO mpls ip
```

To test the configuration:

On BB1

```
BB1#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.67.6 4 msec 0 msec 0 msec
  2 10.1.56.5 [MPLS: Label 509 Exp 0] 96 msec 100 msec 108 msec
  3 10.1.35.3 [MPLS: Label 308 Exp 0] 52 msec 52 msec 56 msec
  4 10.1.23.2 [MPLS: Label 201 Exp 0] 52 msec 56 msec 52 msec
  5 10.1.12.1 28 msec *  28 msec
```

WOW, the customer can see all the labels, let's reload the routers and check again:

On All Routers:

```
#Wr
#Reload

BB1#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.67.6 0 msec 0 msec 0 msec
  2 10.1.56.5 [MPLS: Label 508 Exp 0] 108 msec 100 msec 100 msec
  3 10.1.35.3 [MPLS: Label 306 Exp 0] 56 msec 52 msec 56 msec
  4 10.1.23.2 [MPLS: Label 201 Exp 0] 52 msec 52 msec 56 msec
  5 10.1.12.1 28 msec *  28 msec
```

Note the Traceroute exposes all the links within the provider's network and the label that a given LSR uses to reach the destination.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 13

Reconfigure the appropriate router/s such that a traceroute from BB1 to 1.1.1.1 or R1 to BB1 will not display the links from the provider's network.

The following shows the existing behavior:

On R1

```
R1#Traceroute 7.7.7.7

Type escape sequence to abort.
Tracing the route to 7.7.7.7

  1 10.1.12.2 0 msec 0 msec 0 msec
  2 10.1.23.3 [MPLS: Label 301 Exp 0] 164 msec 100 msec 100 msec
  3 10.1.35.5 [MPLS: Label 505 Exp 0] 56 msec 52 msec 52 msec
  4 10.1.56.6 [MPLS: Label 600 Exp 0] 56 msec 52 msec 52 msec
  5 10.1.67.7 32 msec *  28 msec
```

On BB1

```
BB1#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.67.6 0 msec 4 msec 0 msec
  2 10.1.56.5 [MPLS: Label 508 Exp 0] 104 msec 100 msec 100 msec
  3 10.1.35.3 [MPLS: Label 306 Exp 0] 56 msec 52 msec 56 msec
  4 10.1.23.2 [MPLS: Label 201 Exp 0] 52 msec 52 msec 56 msec
  5 10.1.12.1 28 msec *  28 msec
```

On R2 to R6

The following command has two optional keywords that can be used and they are as follows:

1. **Forwarded: Propagates IP TTL for forwarded traffic, this option will hide the MPLS structure from the customers.**

2. **Local: Propagates IP TTL for locally originated traffic, this option will hide the MPLS structure from the LSRs.**

```
Rx(config)#NO mpls ip propagate-ttl forwarded
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R1

```
R1#Traceroute 7.7.7.7

Type escape sequence to abort.
Tracing the route to 7.7.7.7

  1 10.1.12.2 0 msec 0 msec 0 msec
  2 10.1.56.6 [MPLS: Label 600 Exp 0] 56 msec 52 msec 56 msec
  3 10.1.67.7 28 msec *  28 msec
```

On BB1

```
BB1#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.67.6 0 msec 0 msec 0 msec
  2 10.1.23.2 [MPLS: Label 201 Exp 0] 52 msec 52 msec 56 msec
  3 10.1.12.1 28 msec *  28 msec
```

Note the MPLS structure is hidden from customer's perspective, but the following Traceroute reveals that the LSRs within the cloud will see the structure of the cloud:

On R6

```
R6#Traceroute 1.1.1.1

Type escape sequence to abort.

Tracing the route to 1.1.1.1

  1 10.1.56.5 [MPLS: Label 508 Exp 0] 168 msec 100 msec 100 msec
  2 10.1.35.3 [MPLS: Label 306 Exp 0] 56 msec 52 msec 52 msec
  3 10.1.23.2 [MPLS: Label 201 Exp 0] 56 msec 52 msec 52 msec
  4 10.1.12.1 40 msec *  28 msec
```

On R2

```
R2#Traceroute 7.7.7.7

Type escape sequence to abort.
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Tracing the route to 7.7.7.7

  1 10.1.23.3 [MPLS: Label 301 Exp 0] 176 msec 100 msec 100 msec
  2 10.1.35.5 [MPLS: Label 505 Exp 0] 52 msec 52 msec 56 msec
  3 10.1.56.6 [MPLS: Label 600 Exp 0] 52 msec 52 msec 52 msec
  4 10.1.67.7 32 msec *  28 msec
```

To test the second option (Local) :

On R2 to R6

```
Rx(config)#NO mpls ip propagate-ttl local
```

To verify the configuration:

On BB1

Note the following reveals that the "NO mpls ip propagate-ttl local" command does NOT affect the customers:

```
BB1#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.67.6 0 msec 0 msec 4 msec
  2 10.1.56.5 [MPLS: Label 508 Exp 0] 180 msec 100 msec 100 msec
  3 10.1.35.3 [MPLS: Label 306 Exp 0] 52 msec 52 msec 56 msec
  4 10.1.23.2 [MPLS: Label 201 Exp 0] 52 msec 56 msec 52 msec
  5 10.1.12.1 28 msec *  28 msec
```

On R1

```
R1#Traceroute 7.7.7.7

Type escape sequence to abort.
Tracing the route to 7.7.7.7

  1 10.1.12.2 4 msec 0 msec 0 msec
  2 10.1.23.3 [MPLS: Label 301 Exp 0] 156 msec 100 msec 100 msec
  3 10.1.35.5 [MPLS: Label 505 Exp 0] 52 msec 52 msec 56 msec
  4 10.1.56.6 [MPLS: Label 600 Exp 0] 52 msec 56 msec 52 msec
  5 10.1.67.7 28 msec *  28 msec
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R6

Note the "NO mpls ip propagate-ttl local" command ONLY affected the LSRs:

```
R6#Traceroute 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.23.2 [MPLS: Label 201 Exp 0] 52 msec 56 msec 52 msec
  2 10.1.12.1 28 msec *  28 msec
```

Task 14

Remove the "MPLS IP" command from all interfaces and verify the configuration.

On R2

```
R2#Show mpls interface

Interface              IP            Tunnel   Operational
FastEthernet0/1        Yes (ldp)     No       Yes

R2(config)#int F0/1
R2(config-if)#NO mpls ip
```

To verify the configuration:

On R2

```
R2#Show mpls interface

Interface              IP            Tunnel   Operational
R2#
```

On R3

```
R3#Show mpls interface

Interface              IP            Tunnel   Operational
FastEthernet0/1        Yes (ldp)     No       Yes
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Serial0/0.34            Yes (ldp)      No        Yes
Serial0/0.35            Yes (ldp)      No        Yes

R3(config)#Int S0/0.34
R3(config-subif)#NO Mpls ip

R3(config)#Int S0/0.35
R3(config-subif)#NO Mpls ip

R3(config)#Int F0/1
R3(config-if)#NO Mpls ip
```

To verify the configuration:

On R3

```
R3#Show mpls interface

Interface              IP             Tunnel   Operational
R3#
```

On R4

```
R4#Show mpls interface

Interface              IP             Tunnel   Operational
Serial0/0.43            Yes (ldp)      No        Yes
Serial0/0.45            Yes (ldp)      No        Yes

R4(config)#Int S0/0.43
R4(config-subif)#NO Mpls ip

R4(config)#Int S0/0.45
R4(config-subif)#NO Mpls ip
```

To verify the configuration:

On R4

```
R4#Show mpls interface

Interface              IP             Tunnel   Operational
R4#
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R5

```
R5#Show mpls interface

Interface              IP             Tunnel   Operational
FastEthernet0/1        Yes (ldp)      No       Yes
Serial0/0.53           Yes (ldp)      No       Yes
Serial0/0.54           Yes (ldp)      No       Yes

R5(config)#Int F0/1
R5(config-if)#NO Mpls ip

R5(config)#Int S0/0.53
R5(config-subif)#NO Mpls ip

R5(config)#Int S0/0.54
R5(config-subif)#NO Mpls ip
```

To verify the configuration:

On R5

```
R5#Show mpls interfaces

Interface              IP             Tunnel   Operational
R5#
```

On R6

```
R6#Show mpls interfaces

Interface              IP             Tunnel   Operational
FastEthernet0/1        Yes (ldp)      No       Yes

R6(config)#Int F0/1
R6(config-if)#NO Mpls ip
```

To verify the configuration:

On R6

```
R6#Show mpls interfaces

Interface              IP             Tunnel   Operational
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On BB1

```
BB1#Show mpls interfaces

Interface              IP            Tunnel   Operational
BB1#
```

Task 15

Enable LDP on all the links connecting the routers to each other. DO NOT use the "MPLS IP" command, or a global configuration command to accomplish this task.

LDP is enabled by configuring "MPLS IP" in the interface configuration mode and typically some type of IGP will also need to run on that given interface. To minimize the configuration and errors (Forgetting to enable LDP on a given interface) the LDP Autoconfiguration can be used. When this feature is enabled, LDP is automatically enabled on every interface of that LSR, however, if LDP must be disabled for a given interface, the "NO Mpls ldp igp autoconfig" interface configuration mode can be used. In this case since OSPF is running as the IGP within the MPLS cloud, LDP Autoconfiguration is enabled for all interface/s within area 0.

On BB1

```
BB1(config)#Router ospf 1
BB1(config-router)#Mpls ldp autoconfig area 0
```

To verify the configuration:

On BB1

```
BB1#Show mpls interfaces

Interface              IP            Tunnel   Operational

FastEthernet0/0        Yes (ldp)     No       Yes
FastEthernet0/1        Yes (ldp)     No       Yes
```

To disable LDP on F0/1:

```
BB1(config)#Int F0/1
BB1(config-if)#NO Mpls ldp igp autoconfig
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R6

```
R6(config)#Router ospf 1
R6(config-router)#MPLS ldp autoconfig area 0
```

To verify the configuration:

On R6

```
R6#Show mpls interfaces

Interface              IP            Tunnel   Operational
FastEthernet0/0        Yes (ldp)     No       Yes
FastEthernet0/1        Yes (ldp)     No       Yes
```

You should see the following console message:

```
%LDP-5-NBRCHG: LDP Neighbor 7.7.7.7:0 (1) is UP
```

On R5

```
R5(config)#Router ospf 1
R5(config-router)#Mpls ldp autoconfig area 0
```

To verify the configuration:

On R5

```
R5#Show mpls interfaces

Interface              IP            Tunnel   Operational
FastEthernet0/1        Yes (ldp)     No       Yes
Serial0/0.53           Yes (ldp)     No       Yes
Serial0/0.54           Yes (ldp)     No       Yes
```

You should see the following console message:

```
%LDP-5-NBRCHG: LDP Neighbor 6.6.6.6:0 (1) is UP
```

On R4

```
R4(config)#Router ospf 1
R4(config-router)#Mpls ldp autoconfig area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R4

```
R4#Show mpls interfaces

Interface              IP             Tunnel   Operational
Serial0/0.43           Yes (ldp)      No       Yes
Serial0/0.45           Yes (ldp)      No       Yes
```

You should see the following console message:

```
%LDP-5-NBRCHG: LDP Neighbor 5.5.5.5:0 (1) is UP
```

On R3

```
R3(config)#Router ospf 1
R3(config-router)#Mpls ldp autoconfig area 0
```

To verify the configuration:

On R3

```
R3#Show mpls interfaces

Interface              IP             Tunnel   Operational
FastEthernet0/1        Yes (ldp)      No       Yes
Serial0/0.34           Yes (ldp)      No       Yes
Serial0/0.35           Yes (ldp)      No       Yes
```

You should see the following console message:

```
%LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (1) is UP
%LDP-5-NBRCHG: LDP Neighbor 5.5.5.5:0 (2) is UP
```

On R2

```
R2(config)#Router ospf 1
R2(config-router)#Mpls ldp autoconfig area 0
```

To verify the configuration:

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Show mpls interfaces

Interface              IP             Tunnel   Operational
FastEthernet0/0        Yes (ldp)      No       Yes
FastEthernet0/1        Yes (ldp)      No       Yes
```

You should see the following console message:

```
%LDP-5-NBRCHG: LDP Neighbor 3.3.3.3:0 (1) is UP
```

On R1

```
R1(config)#Router ospf 1
R1(config-router)#Mpls ldp autoconfig area 0
```

To verify the configuration:

On R1

```
R1#Show mpls interfaces

Interface              IP             Tunnel   Operational
FastEthernet0/0        Yes (ldp)      No       Yes
```

You should see the following console message:

```
%LDP-5-NBRCHG: LDP Neighbor 2.2.2.2:0 (1) is UP
```

Task 16

Configure the appropriate router/s such that a failure in one of the links does NOT tear down the LDP session. DO NOT configure a GRE or an IPnIP tunnel to establish this task.

When a link between two LSRs goes down, the two LSRs that share the link will tear down the LDP session, and when the link comes back up, the two LSRs have to re-establish their session and populate the LIB and the LFIB. This bevaior can be changed by protecting the LDP session using a feature called LDP session protection.

When the LDP session protection is configured, a targeted LDP session is built between the routers that have this feature configured. When the directly connected link goes down, the targeted LDP session remains up as long as a redundant link exists between the two LSRs.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

This feature can be enabled using the "MPLS LDP Session Protection" global configuration command.

There are two ways to configure LDP session protection:

1. **Configure the feature on all routers.**
2. **Configure the feature on one router, and then, configure the other router to accept the targeted LDP hellos. This can be accomplished using the "MPLS LDP discovery targeted-hello accept" global configuration command.**

In this case the "MPLS LDP Session Protection" MUST be configured on R3, R4 and R5 because they have redundant links between them.

To see the behavior before configuring this feature:

To see the LDP neighbors before testing the existing condition:

On R3

```
R3#Show mpls ldp neighbor | Inc Peer LDP

    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 3.3.3.3:0
    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 3.3.3.3:0
```

Let's Shutdown the S0/0.34 sub-interface and see the result:

```
R3(config)#Int S0/0.34
R3(config-subif)#Shut
```

You should see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 44.4.4.4 on Serial0/0.34 from FULL to
DOWN, Neighbor Down: Interface down or detached

%LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (3) is DOWN (Interface not
operational)
```

NOTE: Since the link that R3 and R4 share is down, the LDP session is torn down as well. To see the existing LDP sessions:

```
R3#Sh mpls ldp neigh | Inc Peer LDP

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 3.3.3.3:0
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

You can see that R3 has two LDP sessions, one with R2 and another with R5.

Let's bring the S0/0.34 sub-interface back up and configure the "LDP Session protection" feature:

On R3

```
R3(config)#Int S0/0.34
R3(config-subif)#NO Shut
```

You should see the following console messages:

```
%OSPF-5-ADJCHG: Process 1, Nbr 44.4.4.4 on Serial0/0.34 from LOADING to
FULL, Loading Done

%LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (3) is UP
```

To verify the configuration:

On R3

```
R3#Show mpls ldp neighbor | Inc Peer LDP

    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 3.3.3.3:0
    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 3.3.3.3:0
```

On R3, R4 and R5

```
Rx(config)#MPLS LDP Session Protection
```

To verify the configuration:

On R3

```
R3#Show mpls ldp discovery detail

 Local LDP Identifier:
    3.3.3.3:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/1 (ldp): xmit/recv
            Enabled: IGP config;
            Hello interval: 15000 ms; Transport IP addr: 3.3.3.3
            LDP Id: 2.2.2.2:0
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
            Src IP addr: 10.1.23.2; Transport IP addr: 2.2.2.2
            Hold time: 45 sec; Proposed local/peer: 45/45 sec
            Reachable via 2.2.2.2/32
      Serial0/0.34 (ldp): xmit/recv
          Enabled: IGP config;
          Hello interval: 15000 ms; Transport IP addr: 3.3.3.3
          LDP Id: 4.4.4.4:0
            Src IP addr: 10.1.34.4; Transport IP addr: 4.4.4.4
            Hold time: 45 sec; Proposed local/peer: 45/45 sec
            Reachable via 4.4.4.4/32
      Serial0/0.35 (ldp): xmit/recv
          Enabled: IGP config;
          Hello interval: 15000 ms; Transport IP addr: 3.3.3.3
          LDP Id: 5.5.5.5:0
            Src IP addr: 10.1.35.5; Transport IP addr: 5.5.5.5
            Hold time: 45 sec; Proposed local/peer: 45/45 sec
            Reachable via 5.5.5.5/32
  Targeted Hellos:
      3.3.3.3 -> 2.2.2.2 (ldp): active, xmit
          Hello interval: 5000 ms; Transport IP addr: 0.0.0.0
      3.3.3.3 -> 5.5.5.5 (ldp): active/passive, xmit/recv
          Hello interval: 10000 ms; Transport IP addr: 3.3.3.3
          LDP Id: 5.5.5.5:0
            Src IP addr: 5.5.5.5; Transport IP addr: 5.5.5.5
            Hold time: 90 sec; Proposed local/peer: 90/90 sec
            Reachable via 5.5.5.5/32
      3.3.3.3 -> 4.4.4.4 (ldp): active/passive, xmit/recv
          Hello interval: 10000 ms; Transport IP addr: 3.3.3.3
          LDP Id: 4.4.4.4:0
            Src IP addr: 4.4.4.4; Transport IP addr: 4.4.4.4
            Hold time: 90 sec; Proposed local/peer: 90/90 sec
            Reachable via 4.4.4.4/32
```

You can see the targeted Hellos.

To test the configuration:

On R3

```
R3#Debug mpls ldp session protection

R3(config)#Int S0/0.34
R3(config-subif)#Shut
```

You should see the following console messages and Debug output:

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
OSPF adjacency is torn down:
%OSPF-5-ADJCHG: Process 1, Nbr 44.4.4.4 on Serial0/0.34 from FULL to
DOWN, Neighbor Down: Interface down or detached

The Primary LDP adjacency is LOST and the holddown timer is started
counting down:
LDP SP: 4.4.4.4:0: last primary adj lost; starting session protection
holdup timer

By Default, the holdup timer is set for 10 days:
LDP SP: 4.4.4.4:0: LDP session protection holdup timer started, 86400
seconds

The LDP session transitioned from Ready to Protecting:
LDP SP: 4.4.4.4:0: state change (Ready -> Protecting)

The Holdup counter is initiated:
%LDP-5-SP: 4.4.4.4:0: session hold up initiated
```

Lets verify the configuration:

On R3

```
R3#Show mpls ldp neigh | Inc Peer LDP

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 3.3.3.3:0
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 3.3.3.3:0
```

Let's bring up the link and observe the debug and console output:

```
R3(config)#Int S0/0.34
R3(config-subif)#NO Shut
```

You should see the following console messages and debug output:

```
%OSPF-5-ADJCHG: Process 1, Nbr 44.4.4.4 on Serial0/0.34 from LOADING to
FULL, Loading Done

LDP SP: 4.4.4.4:0: primary adj restored; stopping session protection
holdup timer

LDP SP: 4.4.4.4:0: state change (Protecting -> Ready)

%LDP-5-SP: 4.4.4.4:0: session recovery succeeded
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 16

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

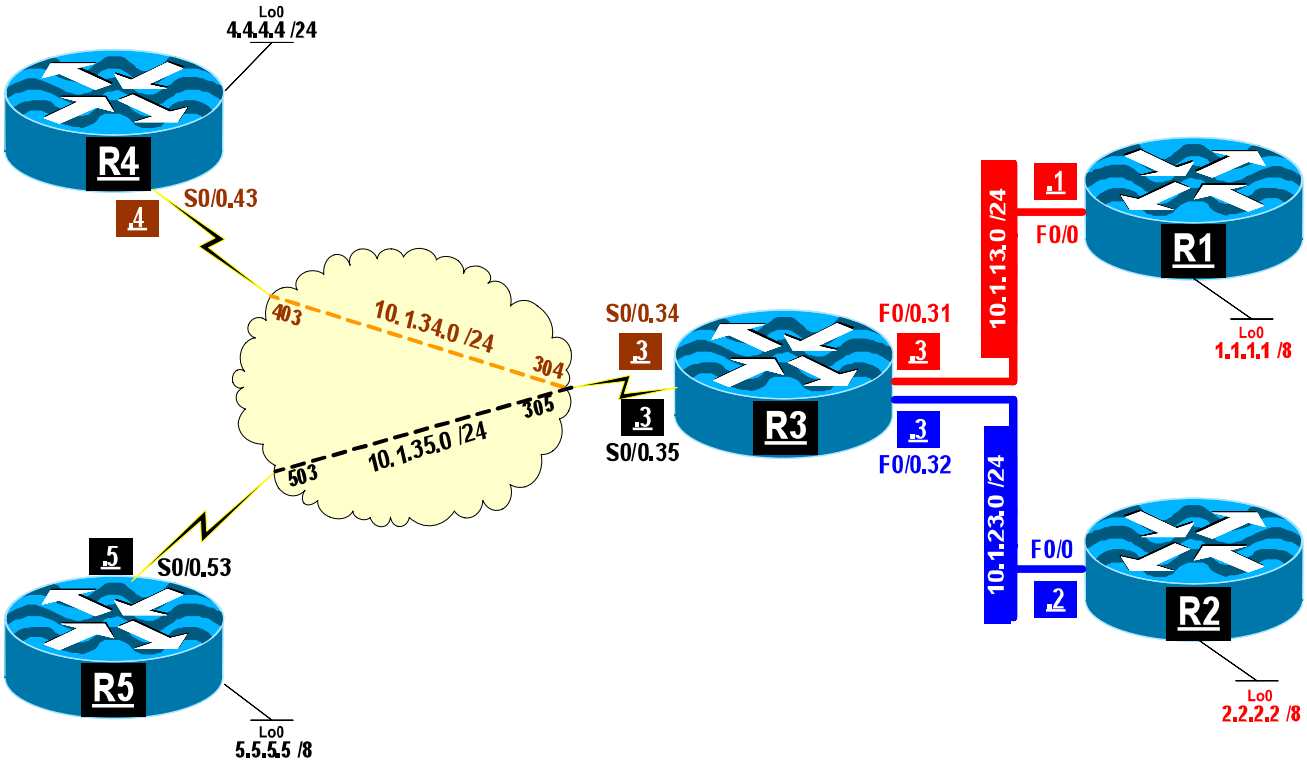
Paul Negron
CCIE #14856
SP

VRF-LITE

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1 – VRF Lite



Lab Setup:

- The frame-relay connection between R3, R4 and R5 should be configured as point-to-point using the DLCIs identified in the diagram.
- R1 should be configured in VLAN 31, whereas, R2 should be configured in VLAN 32.
- R3's F0/0 should be configured as a trunk using an industry standard trunking protocol; R3 should be configured with two sub-interfaces, F0/0.31 and F0/0.32; these sub-interfaces should be connecting to R1 and R2 respectively.
- Use the IP addressing assignment from the following chart:

IP Addressing:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Router	Interface	IP address
R1	Lo0 F0/0	1.1.1.1/ 8 10.1.13.1/ 24
R2	Lo0 F0/0	2.2.2.2 /8 10.1.23.2/ 24
R3	F0/0.31 F0/0.32 S0/0.34 S0/0.35	10.1.13.3 /24 10.1.23.3 /24 10.1.34.3 /24 10.1.35.3/ 24
R4	Lo0 S0/0.43	4.4.4.4 /8 10.1.34.4/ 24
R5	Lo0 S0/0.53	5.5.5.5 /8 10.1.35.5/ 24

Task 1

Configure Eigrp AS 100 on R1 and R4; R1 should advertise its Loopback 0 and F0/0 in this routing domain, whereas, R4 should advertise its Loopback0 and S0/0.43 in this routing domain.

Configure RIPv2 on R2 and R5, R2 should advertise its Loopback0 and F0/0, whereas, R5 should advertise its Loopback0 and S0/0.53 in this routing domain.

On R1

```
R1(config)#Router eigrp 100
R1(config-router)#Netw 1.1.1.1 0.0.0.0
R1(config-router)#Netw 10.1.13.1 0.0.0.0
```

On R2

```
R2(config)#Router rip
R2(config-router)#No au
R2(config-router)#Ver 2
R2(config-router)#Netw 10.0.0.0
R2(config-router)#Netw 2.0.0.0
```

On R4

```
R4(config)#Router eigrp 100
R4(config-router)#No au
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#Netw 4.4.4.4 0.0.0.0
R4(config-router)#Netw 10.1.34.4 0.0.0.0
```

On R5

```
R5(config)#Router rip
R5(config-router)#No au
R5(config-router)#Ver 2
R5(config-router)#Netw 10.0.0.0
R5(config-router)#Netw 5.0.0.0
```

This step prepares the endpoints for routing. The core configuration in the next step is used to bridge connectivity.

Task 2

Configure R3 to provide connectivity to the respective routing domains created in the previous task. R1 and R4 should NOT have any reachability to R2 or R5. You should use VRFs using the following table:

Routers and VRF	RD	RT
R1 and R4 (VRF TSTA)	1:14	1:104
R2 and R5 (VRF TSTB)	1:25	2:205

On R3

```
R3(config)#IP VRF TSTA
R3(config-vrf)#rd  1:14
R3(config-vrf)#route-target both 1:104

R3(config)#IP VRF TSTB
R3(config-vrf)#rd 1:25
R3(config-vrf)#route-target both 1:205

R3(config)#Int F0/0.31
R3(config-subif)#IP VRF forwarding TSTA
```

You should get the following console message:

% Interface FastEthernet0/0.31 IP address 10.1.13.3 removed due to enabling VRF TSTA

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-subif)#IP Address 10.1.13.3 255.255.255.0

R3(config-subif)#Int S0/0.34
R3(config-subif)#IP VRF forwarding TSTA
```

You should get the following console message:

% Interface Serial0/0.34 IP address 10.1.34.3 removed due to enabling VRF TSTA

```
R3(config-subif)#IP Address 10.1.34.3 255.255.255.0

R3(config)#Int F0/0.32
R3(config-subif)#IP VRF forwarding TSTB
```

You should get the following console message:

% Interface FastEthernet0/0.32 IP address 10.1.23.3 removed due to enabling VRF TSTB

```
R3(config-subif)#IP Address 10.1.23.3 255.255.255.0

R3(config)#Int S0/0.35
R3(config-subif)#IP VRF forwarding TSTB
```

You should get the following console message:

% Interface Serial0/0.35 IP address 10.1.35.3 removed due to enabling VRF TSTB

```
R3(config-subif)#IP Address 10.1.35.3 255.255.255.0

R3(config)#Router eigrp 3
R3(config-router)#No Au

R3(config-router)#Address-family IPv4 VRF TSTA
R3(config-router-af)#Autonomous-system 100
R3(config-router-af)#No Au
R3(config-router-af)#Netw 10.1.13.3 0.0.0.0
R3(config-router-af)#Netw 10.1.34.3 0.0.0.0
```

You should get the following console message:

%DUAL-5-NBRCHANGE: IP-EIGRP(1) 100: Neighbor 10.1.34.4 (Serial0/0.34) is up: new adjacency
%DUAL-5-NBRCHANGE: IP-EIGRP(1) 100: Neighbor 10.1.13.1 (FastEthernet0/0.31) is up: new adjacency

```
R3(config)#Router rip
R3(config-router)#No Au
R3(config-router)#Ver 2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#Address-family IPV4 VRF TSTB
R3(config-router-af)#Netw 10.0.0.0
```

To verify and test the Configuration:

On R1

Note: R1 can see the routes from R4 but not R5.

```
R1#Show ip route Eigrp

D    4.0.0.0/8 [90/2300416] via 10.1.13.3, 00:03:39, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
D       10.1.34.0 [90/2172416] via 10.1.13.3, 00:03:39, FastEthernet0/0

R1#Ping 4.4.4.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms

R1#Ping 4.4.4.4 Source lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms
```

The fact that the ping was sent from the loopback as the source shows that there is mutual connectivity.

On R2

```
R2#Show ip route rip

R    5.0.0.0/8 [120/2] via 10.1.23.3, 00:00:03, FastEthernet0/0
     10.0.0.0/24 is subnetted, 2 subnets
R       10.1.35.0 [120/1] via 10.1.23.3, 00:00:03, FastEthernet0/0

R2#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/52/56 ms

R2#Ping 5.5.5.5 Source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/53/56 ms
```

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

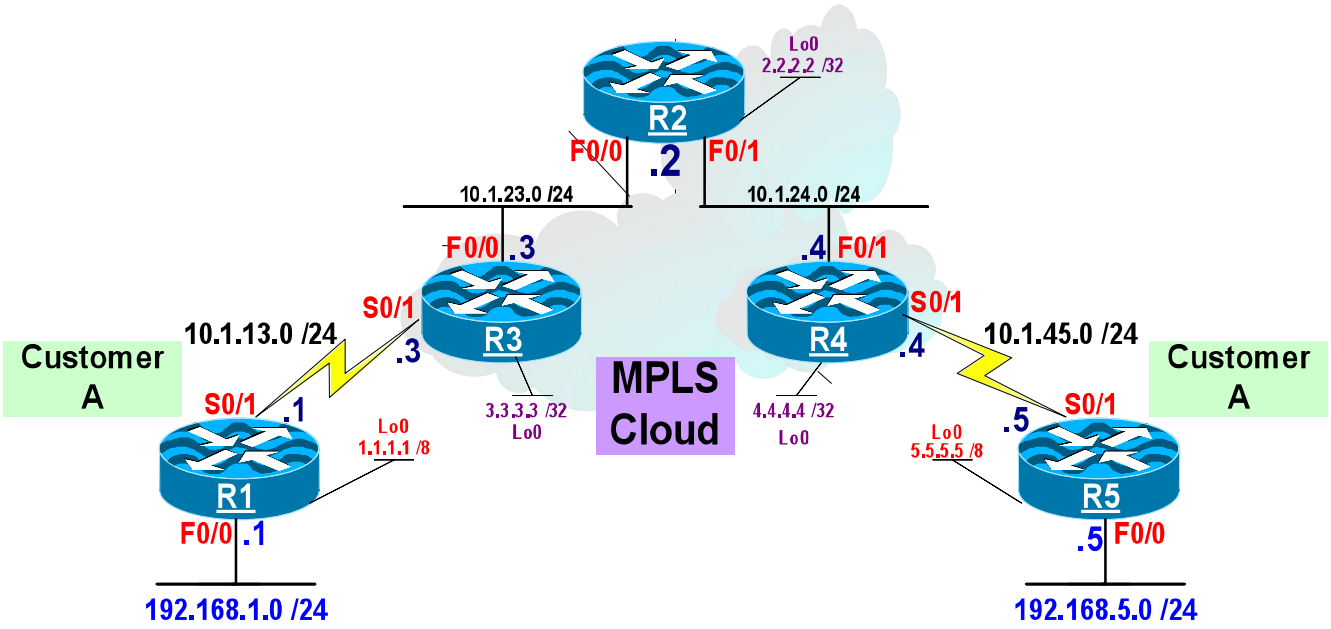
Paul Negron
CCIE #14856
SP

MPLS L3/VPNs

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1
Static & RIPv2 routing in a VPN



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100 and R5 in VLAN 500.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the rest of the routers according to the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.24.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.24.4 0.0.0.0 area 0
```

To verify the configuration:

On R3

```
R3#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       4.4.4.4 [110/3] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       10.1.24.0 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
```

On R2

```
R2#Show ip route ospf | Inc O

O       3.3.3.3 [110/2] via 10.1.23.3, 00:01:03, FastEthernet0/0
O       4.4.4.4 [110/2] via 10.1.24.4, 00:01:03, FastEthernet0/1
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       3.3.3.3 [110/3] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       10.1.23.0 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 200 – 299
R3 – 300 – 399
R4 – 400 – 499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 300 399

R3(config)#int f0/0
R3(config-if)#MPLS IP
```

On R2

```
R2(config)#MPLS label protocol LDP

R2(config)#MPLS LDP router-id lo0

R2(config)#MPLS label range 200 299

R2(config)#int F0/0
R2(config-if)#MPLS IP

R2(config-if)#int F0/1
R2(config-if)#MPLS IP
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#MPLS label protocol LDP

R4(config)#MPLS LDP router-id lo0

R4(config)#MPLS label range 400 499

R4(config)#int F0/1
R4(config-if)#MPLS IP
```

To verify the configuration:

On R4

```
R4#Show mpls interface

Interface              IP            Tunnel   Operational
FastEthernet0/1        Yes (ldp)     No       Yes

R4#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 2.2.2.2.646 - 4.4.4.4.50597
        State: Oper; Msgs sent/rcvd: 10/9; Downstream
        Up time: 00:01:34
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2

R4#Show mpls ldp discovery all

Local LDP Identifier:
    4.4.4.4:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 2.2.2.2:0

R4#Show mpls label range

Downstream Generic label region: Min/Max label: 400/499
```

The default range for the labels is 16 to 100,000 on this platform. These number ranges are software and platform specific. Over 353,000 labels are supported on 6500's as of the writing of this work book.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2#Show mpls interfaces

Interface              IP             Tunnel   Operational
FastEthernet0/0        Yes (ldp)      No       Yes
FastEthernet0/1        Yes (ldp)      No       Yes

R2#Show mpls ldp neighbor

    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 3.3.3.3.23184 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 12/13; Downstream
        Up time: 00:04:01
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.3
        Addresses bound to peer LDP Ident:
          10.1.23.3       10.1.13.3       3.3.3.3
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 4.4.4.4.50597 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 11/13; Downstream
        Up time: 00:03:22
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.4
        Addresses bound to peer LDP Ident:
          10.1.24.4       10.1.45.4       4.4.4.4

R2#Show mpls ldp discovery all

Local LDP Identifier:
    2.2.2.2:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 3.3.3.3:0
        FastEthernet0/1 (ldp): xmit/recv
            LDP Id: 4.4.4.4:0

R2#Show mpls label range

Downstream Generic label region: Min/Max label: 200/299
```

On R3

```
R3#Show mpls interfaces
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
Interface                IP             Tunnel   Operational
FastEthernet0/0          Yes (ldp)      No       Yes

R3#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.23184
        State: Oper; Msgs sent/rcvd: 13/12; Downstream
        Up time: 00:04:01
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2

R3#Show mpls ldp discovery all

Local LDP Identifier:
    3.3.3.3:0
    Discovery Sources:
    Interfaces:
        FastEthernet0/0 (ldp): xmit/recv
            LDP Id: 2.2.2.2:0

R3#Show mpls label range

Downstream Generic label region: Min/Max label: 300/399
```

To verify the LFIB of these routers:

NOTE:

The labels displayed in the output of this lab may differ from your lab but the principle remains the same.

On R3

```
R3#Show mpls forwarding-table

Local  Outgoing    Prefix           Bytes tag   Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched    interface
300    Pop tag     2.2.2.2/32       0           Fa0/0       10.1.23.2
301    201         4.4.4.4/32       0           Fa0/0       10.1.23.2
302    Pop tag     10.1.24.0/24     0           Fa0/0       10.1.23.2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2#Show mpls forwarding-table

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
200    Pop tag     3.3.3.3/32       0          Fa0/0       10.1.23.3
201    Pop tag     4.4.4.4/32       0          Fa0/1       10.1.24.4
```

On R4

```
R4#Show mpls forwarding-table

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
400    Pop tag     2.2.2.2/32       0          Fa0/1       10.1.24.2
401    200         3.3.3.3/32       0          Fa0/1       10.1.24.2
402    Pop tag     10.1.23.0/24     0          Fa0/1       10.1.24.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0
```

Note the exchange of IPv4 routes between BGP neighbors are enabled by default, which means the configured neighbors will NOT only establish a BGP session but they will also receive the advertised prefixes.

Since the "NO bgp default ipv4-unicast" is configured first (Before the neighbor commands), the local router will NOT establish a BGP peer session with any neighbor unless that neighbor is activated.

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#Router bgp 65001
R4(config-router)#NO BGP default ipv4-unicast
R4(config-router)#Neighbor 3.3.3.3 remote-as 65001
R4(config-router)#Neighbor 3.3.3.3 Update-source Lo0
```

To verify the configuration:

On R3

```
R3#Show ip bgp
R3#

R3#Show ip bgp summary
R3#
```

Note there is no IPv4 neighbor adjacency established between the two routers. Since the task states that these two BGP speakers should ONLY establish a VPNv4 peer session, the neighbors MUST be activated in the address-family VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
```

On R4

```
R4(config)#Router bgp 65001
R4(config-router)#Address-family vpnv4 unicast
R4(config-router-af)#Neighbor 3.3.3.3 Activate
```

You should see the following console message:

```
%BGP-5-ADJCHANGE: neighbor 3.3.3.3 Up
```

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4 65001      13      13        1    0    0 00:09:57        0
```

NOTE: The BGP speakers have ONLY established a VPNv4 peering.

On R3

```
R4#Sh ip bgp summ | B Neigh
R4#
```

On R4

```
R4#Sh ip bgp summ | B Neigh
R4#
```

On R2

```
R2#Show ip bgp summary

% BGP not active
```

Note R2 is NOT running BGP at all, and ONLY the edge routers are running BGP. Running BGP on the core routers (P) is NOT necessary; these routers ONLY perform label switching.

Task 4

Configure a **Virtual Routing Forwarding (VRF) Instance with a name of CA** (For Customer A), a **route-distinguisher (RD) of "1:10"** and a **route-target (RT) of "1:100"** on R3. On R4, the same route targets will be used for the vrf, but the **RD should be configured to be "1:20"** and the name of the VRF should be configured as **CB**.

The "IP vrf" command creates a new VRF and enters the global configuration mode for that specific VRF. The name of the VRF is locally significant and it's case sensitive. VRF is NOT operational unless the RD is defined.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The "RD" VRF configuration mode command is used to define and assign an RD to a VRF, remember that the VRF is NOT operational without an RD. RD is a 64 bit value used to transform 32 bit customer IPv4 address, which is NON-Unique into a Unique 96 bit addresses called VPNv4. These addresses are ONLY exchanged between the PE routers and NEVER between the CE routers.

When the CE router sends an update to a PE router, the PE router prepends a 64 bit RD to the IPv4 address (32 bit address) resulting in a globally unique 96 bit address called VPNv4.

The PE router will then send the VPNv4 address/es via MP-BGP session to the other PE router/s. The receiving PE router strips the RD from the VPNv4 prefix, resulting in an IPv4 address.

Remember that the RD does NOT indicate which VRF a given prefix belongs to, it's used to make the VRF prefix/es unique within the cloud. RDs DO NOT identify the VPN.

Let's configure the VRFs and the RDs:

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
```

On R4

```
R4(config)#ip vrf CB
R4(config-vrf)#RD 1:20
```

Since VRFs are locally significant and the RDs will keep the CE routes Unique within the cloud, and the receiving PE will strip off the 64 bit RD value, how does the receiving PE know which VRF does the IP address belong to? The answer is "Route-Target".

The "Route-target import|export" command defines the "RT"; An RT is a BGP extended community that indicates which routes should be exported or imported from MP-BGP into the VRF.

Basically RTs were introduced to support identifying a site that participates in more than one VPN.

The "route-target export" command specifies an RT is to be attached to every route exported from the local VRF to MP-BGP. Whereas, the "route-target import" command specifies an RT to be used as an import filter, ONLY routes matching the RT are imported into the VRF.
This implementation allows a route to have many imported or exported RTs, all to be attached to every imported or exported route.

Let's configure RTs:

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#ip vrf CA
R3(config-vrf)#route-target import 1:100
R3(config-vrf)#route-target export 1:100
```

On R4

```
R4(config)#ip vrf CB
R4(config-vrf)#route-target Both 1:100
```

The "both" keyword is used to replace both "import" and "export" keywords.

But remember that the Rotue-Targets are BGP extended communities that are attached to the VPNv4 addresses, and communities are NOT sent unless they are configured to be sent.

Let's configure the PE routes to send the extended communities:

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

On R4

```
R4(config)#Router bgp 65001
R4(config-router)#Address-family vpnv4 unicast
R4(config-router-af)#Neighbor 3.3.3.3 Send-community extended
```

To verify the VRF Configuration:

On R4

```
R4#Show ip vrf detail

VRF CB; default RD 1:20; default VPNID <not set>
  No interfaces
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  VRF label allocation mode: per-prefix
```

Think of VRFs as VLANs; once a VLAN is configured, the appropriate interface/s are assigned to that VLAN. The same philosophy applies to VRFs, you can think of VRFs as layer three VLANs, in this case, the VRFs are configured, and now the appropriate interface/s needs to be assigned to that VRF:

To associate an interface/s to a given VRF:

On R4

```
R4(config)#int S0/1
R4(config-if)#IP vrf forwarding CB
```

The above command associates an interface with the specified VRF; remember when this command is applied to a given interface, the IP address of that interface is removed and it should be reconfigured.

You should get the following console message:

```
% Interface Serial0/1 IP address 10.1.45.4 removed due to enabling VRF CB

R4(config-if)#ip address 10.1.45.4 255.255.255.0
```

Let's configure R3:

On R3

```
R3(config)#int S0/1
R3(config-if)#ip vrf forwarding CA

% Interface Serial0/1 IP address 10.1.13.3 removed due to enabling VRF CA

R3(config-if)#ip address 10.1.13.3 255.255.255.0
```

To verify the configuration:

On R4

```
R4#Show ip vrf detail

VRF CB; default RD 1:20; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

The above output begins with the name of the VRF, the default Route Distinguisher value and the VPNID which is normally not set by default. The VPNID is an extension used to further identify the VPN by using the customer "OUI" and "vpn index" that can be keyed in as a Hexadecimal or Decimal formatted number.

The interfaces that have the "VRF CB" applied will be listed under the Interfaces: list along with the warning that Connected addresses are no longer in the global routing table. It is a good practice to have at least 1 route target that the router will be both EXPORTING and IMPORTING. The 1:100 Route Target follows this basic design rule.

Currently there is no Export or Import route-maps applied. VRF-SELECT will be discussed later.

On R3

```
R3#Show ip vrf CA

Name                              Default RD          Interfaces
CA                                1:10                Se0/1

R3#Show ip vrf interfaces

Interface              IP-Address      VRF                              Protocol
Se0/1                  10.1.13.3       CA                               up
```

To verify the connectivity between PEs and the CEs:

On R4:

```
R4#Ping 10.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.45.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

A regular ping will no longer work. A ping with no other key words will default to using the global routing table. The 10.1.45.0 /24 prefix is not accessible in the global routing table anymore. The ping must be added with the proper VRF keyword. Remember the IP address is in the VRF and NOT the

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

global routing table.

```
R4#Ping VRF CB 10.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.45.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

On R3:

```
R3#Ping VRF CA 10.1.13.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/29 ms
```

Task 5

Configure a static default route on each Customer router located in CA and CB; these static routes should be configured to point to their respective PE router (**R3** for R1 and **R4** for R5). The PE Routers (R3 and R4), should both be configured with a static route that reaches the loopback and F0/0 interface of the Customer router, R3 and R4 should be able to see both static routes in their BGP table.

To configure a default route on the CEs:

On R1

```
R1(config)#ip route 0.0.0.0 0.0.0.0 10.1.13.3
```

On R5

```
R5(config)#ip route 0.0.0.0 0.0.0.0 10.1.45.4
```

Just a regular default route from the customer's perspective. This method is one of the recommended methods that most Service Providers prefer when offering MPLS VPN as a service to a very small company.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To configure a static route on the PEs:

Note the output of the following show command reveals that when the VRF forwarding was enabled on S0/1 interface, the "Address-family IPv4 VRF CA" was added to the BGP configuration.

On R4

```
R4(config)#ip route vrf CB 5.0.0.0 255.0.0.0 10.1.45.5
R4(config)#ip route vrf CB 192.168.5.0 255.255.255.0 10.1.45.5
```

On R3

```
R3(config)#ip route vrf CA 1.0.0.0 255.0.0.0 10.1.13.1
R3(config)#ip route vrf CA 192.168.1.0 255.255.255.0 10.1.13.1
```

To verify and test the configuration:

On R3

```
R3#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

S    1.0.0.0/8 [1/0] via 10.1.13.1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
S    192.168.1.0/24 [1/0] via 10.1.13.1

R3#Ping vrf CA 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms

R3#Ping vrf CA 192.168.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms
```

On R4

```
R4#Show ip route vrf CB | B Gateway
Gateway of last resort is not set
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
S    5.0.0.0/8 [1/0] via 10.1.45.5
S    192.168.5.0/24 [1/0] via 10.1.45.5
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/1

R4#Ping vrf CB 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

R4#Ping vrf CB 192.168.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

In the last step of this task, the static routes that were configured on the PE routers should be redistributed into BGP so the neighboring PE router can see the route in its BGP table. This redistribution should be done under another address-family called "IPv4 Unicast"; this address-family is created automatically when the VRF, RD and the RTs are configured; the following show command verifies this fact:

On R4

```
R4#Sh run | B router bgp

router bgp 65001
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 3.3.3.3 remote-as 65001
 neighbor 3.3.3.3 update-source Loopback0
 !
 address-family vpnv4
  neighbor 3.3.3.3 activate

  neighbor 3.3.3.3 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf CB
  no synchronization
 exit-address-family
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh run | B router bgp

router bgp 65001
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 4.4.4.4 remote-as 65001
 neighbor 4.4.4.4 update-source Loopback0
 !
 address-family vpnv4
  neighbor 4.4.4.4 activate
  neighbor 4.4.4.4 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf CA
  no synchronization
 exit-address-family
```

To configure the redistribution of the static routes:

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#Address-family IPv4 VRF CA
R3(config-router-af)#Redistribute static
```

Note the redistributed routes should be verified on R4, if the configuration is performed successfully, R4 should be able to see networks 1.0.0.0 /8 and 192.168.1.0 /24 in its vrf CB routing table.

On R4

```
R4#Show ip route vrf CB | B Gateway
Gateway of last resort is not set

B    1.0.0.0/8 [200/0] via 3.3.3.3, 00:02:51
S    5.0.0.0/8 [1/0] via 10.1.45.5
S    192.168.5.0/24 [1/0] via 10.1.45.5
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/1
B    192.168.1.0/24 [200/0] via 3.3.3.3, 00:02:51
```

NOTE: The redistributed routes show up on R4 as BGP routes. Let's configure R4 to redistribute its static routes:

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#Router bgp 65001
R4(config-router)#Address-family IPv4 vrf CB
R4(config-router-af)#Redistribute static
```

To verify the configuration:

On R3

```
R3#Show ip route vrf CA | b Gateway
Gateway of last resort is not set

S    1.0.0.0/8 [1/0] via 10.1.13.1
B    5.0.0.0/8 [200/0] via 4.4.4.4, 00:00:27
B    192.168.5.0/24 [200/0] via 4.4.4.4, 00:00:27
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
S    192.168.1.0/24 [1/0] via 10.1.13.1
```

To verify the routes in VPNv4:

On R3

```
R3#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                0         32768 ?
*>i5.0.0.0          4.4.4.4                  0    100      0 ?
*> 192.168.1.0      10.1.13.1                0         32768 ?
*>i192.168.5.0      4.4.4.4                  0    100      0 ?
Route Distinguisher: 1:20
*>i5.0.0.0          4.4.4.4                  0    100      0 ?
*>i192.168.5.0      4.4.4.4                  0    100      0 ?
```

On R4

```
R4#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10
*>i1.0.0.0          3.3.3.3                  0    100      0 ?
*>i192.168.1.0      3.3.3.3                  0    100      0 ?
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0          3.3.3.3                  0    100      0 ?
*> 5.0.0.0          10.1.45.5                0         32768 ?
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*>i192.168.1.0      3.3.3.3                  0    100      0 ?
*> 192.168.5.0      10.1.45.5                0         32768 ?
```

One of the reasons to use different Route Distinguishers for ALL VRFs in the network is for clarity, especially when trying to figure out which VPN a route belongs to and where it's originated. As you can see, the "1:20" RD is local, but the 1.0.0.0 & 192.168.1.0 networks are present in both RDs, which means that 1.0.0.0 & 192.168.1.0 networks are participating in the same VPN. Even though the VPN is defined by the Route Target, it is still apparent that the prefixes are being imported and exported into both VRFs.

To test the configuration:

On R1

```
R1#Ping 5.5.5.5 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R1#Ping 5.5.5.5 source 192.168.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 84/84/88 ms
```

On R5

```
R5#Ping 1.1.1.1 source 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 5.5.5.5
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/57 ms

R5#Ping 1.1.1.1 source 192.168.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.5.5
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 84/84/84 ms
```

Note the source of the two ping commands are specified to be the loopback 0 & F0/0's IP address, if the source is NOT specified, the Ping will NOT be successful:

On R5

```
R5#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R1

```
R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The ping failed because the source was NOT set based on the loopback0 or the F0/0 interface, if the source is not specified, the source IP address will be the closest interface, in this case the IP address of the S0/1 interface. In order to provide the capability of Pinging without specifying the source IP address, the PE routers should also redistribute the IP address of their S0/1 interface, this can be done using the "Redistribute Connected":

On R3

```
R3(config)#router bgp 65001
R3(config-router)#Address-family IPv4 vrf CA
R3(config-router-af)#Redistribute connected
```

To verify the configuration:

On R4

```
R4#Show ip route vrf CB | B Gateway
Gateway of last resort is not set

B    1.0.0.0/8 [200/0] via 3.3.3.3, 00:21:41
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
S    5.0.0.0/8 [1/0] via 10.1.45.5
S    192.168.5.0/24 [1/0] via 10.1.45.5
     10.0.0.0/24 is subnetted, 2 subnets
B       10.1.13.0 [200/0] via 3.3.3.3, 00:00:26
C       10.1.45.0 is directly connected, Serial0/1
B    192.168.1.0/24 [200/0] via 3.3.3.3, 00:21:41
```

On R4

```
R4(config)#router bgp 65001
R4(config-router)#Address-family IPv4 vrf CB
R4(config-router-af)#Redistribute connected
```

To verify the configuration:

On R3

```
R3#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

S    1.0.0.0/8 [1/0] via 10.1.13.1
B    5.0.0.0/8 [200/0] via 4.4.4.4, 00:17:30
B    192.168.5.0/24 [200/0] via 4.4.4.4, 00:17:30
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.13.0 is directly connected, Serial0/1
B       10.1.45.0 [200/0] via 4.4.4.4, 00:00:45
S    192.168.1.0/24 [1/0] via 10.1.13.1
```

To test the connectivity:

On R1

```
R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

On R5

```
R5#Ping 1.1.1.1

Type escape sequence to abort.
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/57 ms
```

To verify the extended community attached to a given prefix:

On R4

```
R4#Show ip bgp VPNv4 all 5.0.0.0

BGP routing table entry for 1:20:5.0.0.0/8, version 8
Paths: (1 available, best #1, table CB)
  Advertised to update-groups:
        1
 Local
  10.1.45.5 from 0.0.0.0 (4.4.4.4)
   Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced,
best
      Extended Community: RT:1:100
      mpls labels in/out 403/nolabel
```

We have successfully verified that the route-target has been added to the routes as they are exported from the local VRF. This is important as NO vrf will be able to import this route if the extended community is not added to the prefix. The VPNv4 label has also been added to the prefix. The "403(in) and the nolabel(out)" shows that this router expects to see this route with label 403 when it recieves the data traffic and will remove the label before forwarding it to the Customer (R5) router.

```
R4#Show ip bgp VPNv4 all labels

   Network          Next Hop      In label/Out label
Route Distinguisher: 1:10
   1.0.0.0          3.3.3.3         nolabel/303
   10.1.13.0/24     3.3.3.3         nolabel/305
   192.168.1.0      3.3.3.3         nolabel/304
Route Distinguisher: 1:20 (CB)
   1.0.0.0          3.3.3.3         nolabel/303
   5.0.0.0          10.1.45.5       403/nolabel
   10.1.13.0/24     3.3.3.3         nolabel/305
   10.1.45.0/24     0.0.0.0         405/aggregate(CB)
   192.168.1.0      3.3.3.3         nolabel/304
   192.168.5.0      10.1.45.5       404/nolabel
```

This command displays the VPNv4 labels that are added to the routes. A common question that is often asked is "What does the aggregate keyword mean?" The aggregate keyword means that the route requires an IP lookup to determine the next hop of this packet. When we redistributed the connected routes, the far end does not know about this next hop due to the fact that this route is not in the VPN.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

By redistributing the connected route, a normal ping from the interface closest to the destination can be performed by R1 due to the fact that R5 now knows about the source of the packet.

```
R4#Sh mpls forwarding-table vrf CB
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
403	Untagged	5.0.0.0/8[V]	2080	Se0/1	point2point
404	Untagged	192.168.5.0/24[V]	520	Se0/1	point2point
405	Aggregate	10.1.45.0/24[V]	520		

On R3

```
R3#Show mpls forwarding-table vrf CA
```

Local tag	Outgoing tag or VC	Prefix or Tunnel Id	Bytes tag switched	Outgoing interface	Next Hop
303	Untagged	1.0.0.0/8[V]	2600	Se0/1	point2point
304	Untagged	192.168.1.0/24[V]	520	Se0/1	point2point
305	Aggregate	10.1.13.0/24[V]	520		

Task 6

Remove the Static routes and replace the current method of routing between the PE and Customers with RIPv2 Routing.

Removing the Static Configuration on the CE routers:

On R1

```
R1(config)#NO ip route 0.0.0.0 0.0.0.0
```

On R5

```
R5(config)#NO ip route 0.0.0.0 0.0.0.0
```

Removing the Static Configuration on the PE routers:

On R3

```
R3(config)#NO ip route vrf CA 1.0.0.0 255.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#NO ip route vrf CA 192.168.1.0 255.255.255.0

R3(config)#router bgp 65001
R3(config-router)# address ipv4 vrf CA
R3(config-router-af)#NO redistribute static
R3(config-router-af)#NO redistribute connected
```

On R4

```
R4(config)#NO ip route vrf CB 5.0.0.0 255.0.0.0
R4(config)#NO ip route vrf CB 192.168.5.0 255.255.255.0

R4(config)#router bgp 65001
R4(config-router)#address ipv4 vrf CB
R4(config-router-af)#NO redistribute static
R4(config-router-af)#NO redistribute connected
```

Configuring RIPv2 routing on CE routers:

On R1

```
R1(config)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#netw 10.0.0.0
R1(config-router)#netw 192.168.1.0
R1(config-router)#netw 1.0.0.0
```

On R5

```
R5(config)#router rip
R5(config-router)#No au
R5(config-router)#ver 2
R5(config-router)#network 10.0.0.0

R5(config-router)#netw 5.0.0.0
R5(config-router)#netw 192.168.5.0
```

Configuring RIPv2 routing on PE routers:

On R3

```
R3(config)#router rip
R3(config-router)#ver 2
R3(config-router)#Address-family ipv4 vrf CA
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router-af)#Network 10.0.0.0
R3(config-router-af)#version 2
R3(config-router-af)#No au
```

On R4

```
R4(config)#router rip
R4(config-router)#ver 2

R4(config-router)#Address-family ipv4 vrf CB
R4(config-router-af)#Network 10.0.0.0
R4(config-router-af)#No au
R4(config-router-af)#ver 2
```

To verify the configuration:

On R3

```
R3#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

R    1.0.0.0/8 [120/1] via 10.1.13.1, 00:00:20, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
R    192.168.1.0/24 [120/1] via 10.1.13.1, 00:00:20, Serial0/1
```

On R4

```
R4#Show ip route vrf CB | B Gateway
Gateway of last resort is not set

R    5.0.0.0/8 [120/1] via 10.1.45.5, 00:00:23, Serial0/1
R    192.168.5.0/24 [120/1] via 10.1.45.5, 00:00:23, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/1
```

Note the routes are in the appropriate VRFs, the next step is to redistribute the routes.

The redistribution of RIPv2 into MP-BGP is necessary for the routes to show up on the other PE Router.

The redistribution of MP-BGP into RIP is necessary for the local router to translate the routes that have been received by the remote PE router so the CE router/s can see the routes.

Step One

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

RIP routes are redistributed into the BGP for the specified VRF:

On R4

```
R4(config)#router bgp 65001
R4(config-router)#Address-family ipv4 vrf CB
R4(config-router-af)#Redistribute RIP
```

To verify the configuration:

Because of the redistribution, the PE router (R3) on the other side can now see the routes in its BGP table:

On R3

```
R3#Show ip bgp vpnv4 vrf CA | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*>i5.0.0.0          4.4.4.4                  1    100      0 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
*>i192.168.5.0      4.4.4.4                  1    100      0 ?
```

Step Two

In this step the routes are redistributed into RIP, so the CE router will have them in its routing table:

On R3

When the routes are redistributed into RIP, a metric of 5 is assigned; this is done on purpose so they can easily be identified in the routing table of the Customer router (R1):

```
R3(config)#router rip
R3(config-router)#Address-family ipv4 vrf CA
R3(config-router-af)#redistribute BGP 65001 metric 5
```

To verify the configuration:

On R1

Note R1 (The CE router) has the routes in its routing table:

```
R1#Show ip route rip | Inc R
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R    5.0.0.0/8 [120/5] via 10.1.13.3, 00:00:08, Serial0/1
R    192.168.5.0/24 [120/5] via 10.1.13.3, 00:00:08, Serial0/1
R       10.1.45.0 [120/5] via 10.1.13.3, 00:00:08, Serial0/1
```

The same needs to be done in reverse order from R3 to R4 and verified on R5:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#Address-family ipv4 vrf CA
R3(config-router-af)#Redistribute RIP
```

On R4

```
R4(config)#Router rip
R4(config-router)#Address-family ipv4 vrf CB
R4(config-router-af)#Redistribute BGP 65001 metric 5
```

To verify the configuration:

On R5

```
R5#Show ip route rip | Inc R

R    1.0.0.0/8 [120/5] via 10.1.45.4, 00:00:05, Serial0/1
R       10.1.13.0 [120/5] via 10.1.45.4, 00:00:05, Serial0/1
R    192.168.1.0/24 [120/5] via 10.1.45.4, 00:00:05, Serial0/1
```

To see the full picture, the following show commands verifies network 1.0.0.0 /8 from R1 all the way to the upstream router R5; this is called the control plane.

On R1

```
R1#Show ip route 1.0.0.0

Routing entry for 1.0.0.0/8
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via rip
  Advertised by rip
  Routing Descriptor Blocks:
  * directly connected, via Loopback0
      Route metric is 0, traffic share count is 1

R1#Show ip rip database 1.0.0.0 255.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
1.0.0.0/8    directly connected, Loopback0
```

On R3

```
R3#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

R    1.0.0.0/8 [120/1] via 10.1.13.1, 00:00:04, Serial0/1
B    5.0.0.0/8 [200/1] via 4.4.4.4, 00:29:36
B    192.168.5.0/24 [200/1] via 4.4.4.4, 00:29:36
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.13.0 is directly connected, Serial0/1
B       10.1.45.0 [200/0] via 4.4.4.4, 00:29:36
R    192.168.1.0/24 [120/1] via 10.1.13.1, 00:00:04, Serial0/1

R3#Show ip route vrf CA 1.0.0.0

Routing entry for 1.0.0.0/8
  Known via "rip", distance 120, metric 1
  Redistributing via bgp 65001, rip
  Advertised by bgp 65001
  Last update from 10.1.13.1 on Serial0/1, 00:00:11 ago
  Routing Descriptor Blocks:
  * 10.1.13.1, from 10.1.13.1, 00:00:11 ago, via Serial0/1
      Route metric is 1, traffic share count is 1

R3#Show ip bgp vpnv4 vrf CA | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                1         32768 ?

*>i5.0.0.0          4.4.4.4                  1    100      0 ?
*> 10.1.13.0/24     0.0.0.0                  0         32768 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
*> 192.168.1.0      10.1.13.1                1         32768 ?
*>i192.168.5.0      4.4.4.4                  1    100      0 ?

R3#Show ip bgp VPNV4 vrf CA 1.0.0.0/8

BGP routing table entry for 1:10:1.0.0.0/8, version 32
Paths: (1 available, best #1, table CA)
  Advertised to update-groups:
        1
  Local
    10.1.13.1 from 0.0.0.0 (3.3.3.3)
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
      Origin incomplete, metric 1, localpref 100, weight 32768, valid, sourced,
best
      Extended Community: RT:1:100
      mpls labels in/out 303/nolabel
```

Note the Origin code is (incomplete) because the route was redistributed and therefore, the origin of the route is unknown. The Weight attribute is 32768, because the local router is advertising the route in BGP. The extended community is identified, this is the configured route-target, and the last line states that if the local router receives a packet with a label of 303 it will remove the label and forward it.

On R4

```
R4#Show ip bgp vpnv4 vrf CB 1.0.0.0

BGP routing table entry for 1:20:1.0.0.0/8, version 32
Paths: (1 available, best #1, table CB)
  Not advertised to any peer
  Local, imported path from 1:10:1.0.0.0/8
    3.3.3.3 (metric 3) from 3.3.3.3 (3.3.3.3)
      Origin incomplete, metric 1, localpref 100, valid, internal, best
      Extended Community: RT:1:100
      mpls labels in/out nolabel/303

R4#Show ip bgp vpnv4 vrf CB | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0          3.3.3.3                  1    100      0 ?
*> 5.0.0.0          10.1.45.5                1         32768 ?
*>i10.1.13.0/24     3.3.3.3                  0    100      0 ?
*> 10.1.45.0/24     0.0.0.0                  0         32768 ?
*>i192.168.1.0      3.3.3.3                  1    100      0 ?
*> 192.168.5.0      10.1.45.5                1         32768 ?

R4#Show ip route vrf CB | B Gateway
Gateway of last resort is not set

B    1.0.0.0/8 [200/1] via 3.3.3.3, 00:11:50
R    5.0.0.0/8 [120/1] via 10.1.45.5, 00:00:05, Serial0/1
R    192.168.5.0/24 [120/1] via 10.1.45.5, 00:00:05, Serial0/1
     10.0.0.0/24 is subnetted, 2 subnets
B       10.1.13.0 [200/0] via 3.3.3.3, 00:11:50
C       10.1.45.0 is directly connected, Serial0/1
B    192.168.1.0/24 [200/1] via 3.3.3.3, 00:11:50

R4#Sh ip route vrf CB 1.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Routing entry for 1.0.0.0/8
  Known via "bgp 65001", distance 200, metric 1, type internal
  Redistributing via rip
  Advertised by rip metric 5
  Last update from 3.3.3.3 00:13:11 ago
  Routing Descriptor Blocks:
  * 3.3.3.3 (Default-IP-Routing-Table), from 3.3.3.3, 00:13:11 ago
      Route metric is 1, traffic share count is 1
      AS Hops 0
```

On R5

```
R5#Show ip route rip

R    1.0.0.0/8 [120/5] via 10.1.45.4, 00:00:04, Serial0/1
R       10.1.13.0 [120/5] via 10.1.45.4, 00:00:04, Serial0/1
R    192.168.1.0/24 [120/5] via 10.1.45.4, 00:00:04, Serial0/1
```

To test the configuration:

On R5

```
R5#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/57 ms
```

On R1

```
R1#Ping 5.5.5.5

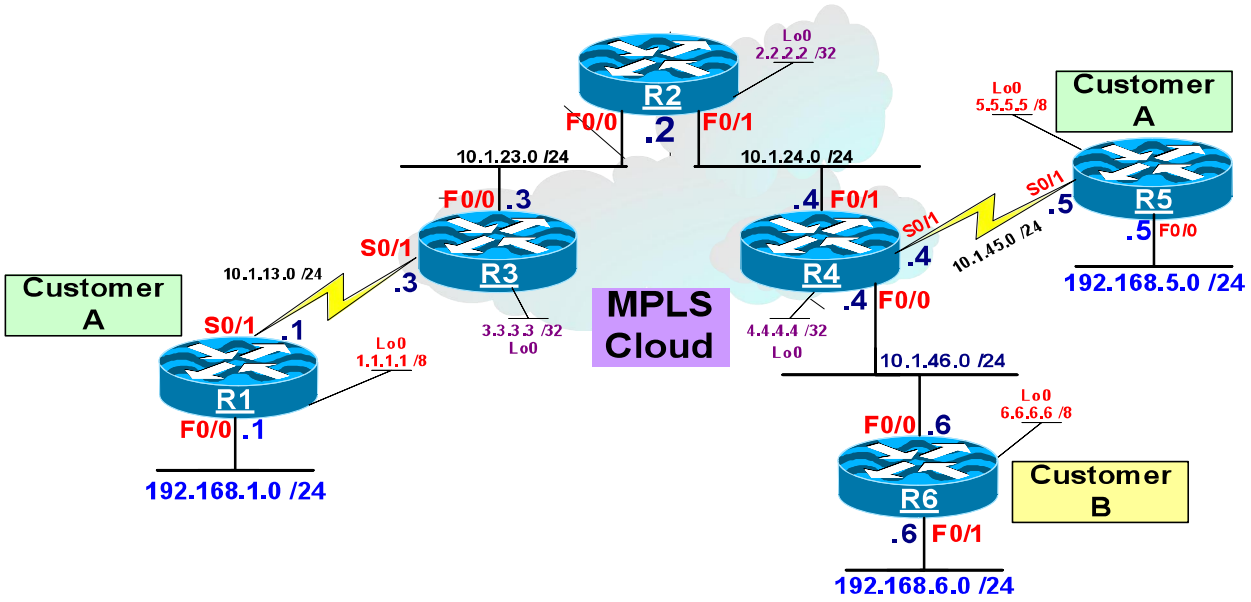
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 2
OSPF Routing in a VPN



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100, R5 in VLAN 500 and F0/1 interface of R6 in VLAN 600.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the F0/0 interface of R4 and R6 in VLAN 900.
- Configure the rest of the routers according to the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.24.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.24.4 0.0.0.0 area 0
```

To verify the configuration:

On R3

```
R3#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       4.4.4.4 [110/3] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       10.1.24.0 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
```

On R2

```
R2#Show ip route ospf | Inc O

O       3.3.3.3 [110/2] via 10.1.23.3, 00:01:03, FastEthernet0/0
O       4.4.4.4 [110/2] via 10.1.24.4, 00:01:03, FastEthernet0/1
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Show ip route ospf | Inc O

O        2.2.2.2 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
O        3.3.3.3 [110/3] via 10.1.24.2, 00:01:41, FastEthernet0/1
O        10.1.23.0 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 200 – 299
R3 – 300 – 399
R4 – 400 – 499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 300 399

R3(config)#int f0/0
R3(config-if)#MPLS IP
```

On R2

```
R2(config)#MPLS label protocol LDP

R2(config)#MPLS LDP router-id lo0

R2(config)#MPLS label range 200 299

R2(config)#int F0/0
R2(config-if)#MPLS IP

R2(config-if)#int F0/1
R2(config-if)#MPLS IP
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#MPLS label protocol LDP

R4(config)#MPLS LDP router-id lo0

R4(config)#MPLS label range 400 499

R4(config)#int F0/1
R4(config-if)#MPLS IP
```

To verify the configuration:

On R4

```
R4#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 2.2.2.2.646 - 4.4.4.4.50597
        State: Oper; Msgs sent/rcvd: 10/9; Downstream
        Up time: 00:01:34
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

On R2

```
R2#Show mpls ldp neighbor

    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 3.3.3.3.23184 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 12/13; Downstream
        Up time: 00:04:01
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.3
        Addresses bound to peer LDP Ident:
          10.1.23.3       10.1.13.3       3.3.3.3
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 4.4.4.4.50597 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 11/13; Downstream
        Up time: 00:03:22
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.4
        Addresses bound to peer LDP Ident:
          10.1.24.4       10.1.45.4       4.4.4.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.23184
        State: Oper; Msgs sent/rcvd: 13/12; Downstream
        Up time: 00:04:01
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 send-community extended
```

On R4

```
R4(config)#Router bgp 65001
R4(config-router)#NO BGP default ipv4-unicast
R4(config-router)#Neighbor 3.3.3.3 remote-as 65001
R4(config-router)#Neighbor 3.3.3.3 Update-source Lo0


R4(config-router)#Address-family vpnv4 unicast
R4(config-router-af)#Neighbor 3.3.3.3 Activate
R4(config-router-af)#Neighbor 3.3.3.3 send-community extended
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

```
R4#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4 65001      13      13        1    0    0 00:09:57        0
```

Task 4

Configure VRFs on R3 and R4 and enable VRF forwarding on the interfaces of these two routers based on the following chart.

Router	VRF	RD	RT	Interface
R3	CA	1:10	1:105	S0/1
R4	CA	1:50	1:105	S0/1
	CB	1:60	1:600	F0/0

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
R3(config-vrf)#route-target both 1:105

R3(config-vrf)#int S0/1
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip addr 10.1.13.3 255.255.255.0
```

On R4

```
R4(config)#ip vrf CA
R4(config-vrf)#rd 1:50
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-vrf)#route-target both 1:105

R4(config-vrf)#int S0/1
R4(config-if)#ip vrf forwarding CA
R4(config-if)#ip addr 10.1.45.4 255.255.255.0

R4(config)#ip vrf CB
R4(config-vrf)#rd 1:60
R4(config-vrf)#route-target both 1:600

R4(config)#int F0/0
R4(config-if)#ip vrf forwarding CB
R4(config-if)#ip addr 10.1.46.4 255.255.255.0
```

To verify the configuration:

On R3

```
R3#Sh ip vrf det

VRF CA; default RD 1:10; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:105
  Import VPN route-target communities
    RT:1:105
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

On R4

```
R4#Show ip vrf det | S CA

VRF CA; default RD 1:50; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:105
  Import VPN route-target communities
    RT:1:105
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix

R4#Show ip vrf det | S CB

VRF CB; default RD 1:60; default VPNID <not set>
Interfaces:
    Fa0/0
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:600
  Import VPN route-target communities
    RT:1:600
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Task 5

Configure Customers R1 and R5 with a VRF service that incorporates OSPF as the routing protocol; R3 should use OSPF process-id of 3, whereas, R4 should use process-id of 4 for this purpose, R1 and R5's loopback0 & F0/0 interfaces should be advertised with their correct mask.
Configure the VRF applied to R6 to use RIPv2 and ensure that this customer participates in the same Closed User Group (CUG) as R1 and R5. R6 should advertise its Loopback0 and F0/0 interfaces in its routing domain.

To configure the CE routers:

On R1

```
R1(config)#router ospf 1

R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#Netw 10.1.13.1 0.0.0.0 area 0
R1(config-router)#Netw 192.168.1.1 0.0.0.0 area 0

R1(config-router)#Int lo0
R1(config-if)#ip ospf netw point-to-point
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R5

```
R5(config)#router ospf 1
R5(config-router)#Netw 5.5.5.5 0.0.0.0 area 0
R5(config-router)#Netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#Netw 192.168.5.5 0.0.0.0 area 0

R5(config-router)#int lo0
R5(config-if)#ip ospf net point-to-point
```

On R6

```
R6(config)#router rip
R6(config-router)#no au
R6(config-router)#ver 2
R6(config-router)#Netw 192.168.6.0
R6(config-router)#Netw 6.0.0.0
R6(config-router)#Netw 10.0.0.0
```

To configure the PE routers:

On R3

```
R3(config)#Router ospf 3 vrf CA
R3(config-router)#network 10.1.13.3 0.0.0.0 area 0
```

On R4

```
R4(config)#router ospf 4 vrf CA
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 0

R4(config)#router rip
R4(config-router)#no au
R4(config-router)#ver 2
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#Network 10.0.0.0
R4(config-router-af)#no au

R4(config-router-af)#ver 2
```

To verify the configuration:

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh ip ospf nei

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           1   FULL/BDR        00:00:35    10.1.23.2       FastEthernet0/0
1.1.1.1           0   FULL/  -        00:00:31    10.1.13.1       Serial0/1
```

1.1.1.1 is R1 (The CE) and 2.2.2.2 is the OSPF adjacency with R2 (The P router).

```
R3#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

O    1.0.0.0/8 [110/65] via 10.1.13.1, 00:06:06, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
O    192.168.1.0/24 [110/65] via 10.1.13.1, 00:06:06, Serial0/1
```

On R4

```
R4#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           1   FULL/DR         00:00:36    10.1.24.2       FastEthernet0/1
5.5.5.5           0   FULL/  -        00:00:31    10.1.45.5       Serial0/1

R4#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

O    5.0.0.0/8 [110/65] via 10.1.45.5, 00:09:50, Serial0/1
O    192.168.5.0/24 [110/65] via 10.1.45.5, 00:00:23, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/1

R4#Show ip route vrf CB | B Gate
Gateway of last resort is not set

O    5.0.0.0/8 [110/65] via 10.1.45.5, 00:07:36, Serial0/1
O    192.168.5.0/24 [110/65] via 10.1.45.5, 00:07:36, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/1
```

The CE routes are redistribted into MP-BGP:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router-af)#redistribute OSPF 3
```

On R4

```
R4(config)#router bgp 65001
R4(config-router)#address-family ipv4 vrf CA
R4(config-router-af)#redistribute OSPF 4

R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#redistribute rip
```

To verify the configuration:

On R3

```
R3#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                65         32768 ?
*>i5.0.0.0          4.4.4.4                  65    100      0 ?
*> 10.1.13.0/24     0.0.0.0                   0         32768 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
*> 192.168.1.0      10.1.13.1                65         32768 ?
*>i192.168.5.0      4.4.4.4                  65    100      0 ?
Route Distinguisher: 1:50
*>i5.0.0.0          4.4.4.4                  65    100      0 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
*>i192.168.5.0      4.4.4.4                  65    100      0 ?
```

Since none of the VRFs configured on R3 have the "route-target import 1:600", R3 will NOT have the routes from R6 in its table. The following configuration resolves this problem:

On R4

```
R4(config)#ip vrf CB
R4(config-vrf)#route-target both 1:615

R4(config)#ip vrf CA
R4(config-vrf)#route-target both 1:615
```

On R3

```
R3(config)#ip vrf CA
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-vrf)#route-target both 1:615
```

To verify the configuration:

On R3

```
R3#Show ip bgp vpnv4 all | B Net
```

```
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                65         32768 ?
*>i5.0.0.0          4.4.4.4                  65    100      0 ?
*>i6.0.0.0          4.4.4.4                   1    100      0 ?
*> 10.1.13.0/24     0.0.0.0                   0         32768 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
*>i10.1.46.0/24     4.4.4.4                   0    100      0 ?
*> 192.168.1.0      10.1.13.1                65         32768 ?
*>i192.168.5.0      4.4.4.4                  65    100      0 ?
*>i192.168.6.0      4.4.4.4                   1    100      0 ?
Route Distinguisher: 1:50
*>i5.0.0.0          4.4.4.4                  65    100      0 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
*>i192.168.5.0      4.4.4.4                  65    100      0 ?
Route Distinguisher: 1:60
*>i6.0.0.0          4.4.4.4                   1    100      0 ?
*>i10.1.46.0/24     4.4.4.4                   0    100      0 ?
*>i192.168.6.0      4.4.4.4                   1    100      0 ?
```

On R4

```
R4#Show ip bgp vpnv4 all | B Network
```

```
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10
*>i1.0.0.0          3.3.3.3                  65    100      0 ?
*>i10.1.13.0/24     3.3.3.3                   0    100      0 ?
*>i192.168.1.0      3.3.3.3                  65    100      0 ?
Route Distinguisher: 1:50 (default for vrf CA)
*>i1.0.0.0          3.3.3.3                  65    100      0 ?
*> 5.0.0.0          10.1.45.5                65         32768 ?
*> 6.0.0.0          10.1.46.6                 1         32768 ?
*>i10.1.13.0/24     3.3.3.3                   0    100      0 ?
*> 10.1.45.0/24     0.0.0.0                   0         32768 ?
*> 10.1.46.0/24     0.0.0.0                   0         32768 ?
*>i192.168.1.0      3.3.3.3                  65    100      0 ?
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 192.168.5.0       10.1.45.5                65         32768 ?
*> 192.168.6.0       10.1.46.6                 1         32768 ?
Route Distinguisher: 1:60 (default for vrf CB)
*>i1.0.0.0           3.3.3.3                  65    100      0 ?
*> 5.0.0.0           10.1.45.5                65         32768 ?
*> 6.0.0.0           10.1.46.6                 1         32768 ?
*>i10.1.13.0/24      3.3.3.3                   0    100      0 ?
*> 10.1.45.0/24      0.0.0.0                   0         32768 ?
*> 10.1.46.0/24      0.0.0.0                   0         32768 ?
*>i192.168.1.0       3.3.3.3                  65    100      0 ?
*> 192.168.5.0       10.1.45.5                65         32768 ?
*> 192.168.6.0       10.1.46.6                 1         32768 ?
```

Redistribute BGP into OSPF on R3, and into OSPF and RIP on R4 so the CE routers (R1, R5 and R6) can see the routes in their routing table:

On R3

```
R3(config)#router ospf 3 vrf CA
R3(config-router)#redistribute bgp 65001 subnets
```

On R4

```
R4(config)#router ospf 4 vrf CA
R4(config-router)#redistribute BGP 65001 subnets

R4(config)#Router rip
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#redistribute BGP 65001 metric 5
```

To verify the configuration:

On R1

```
R1#Show ip route ospf | Inc O

O E2 5.0.0.0/8 [110/65] via 10.1.13.3, 00:06:23, Serial0/1
O E2 6.0.0.0/8 [110/1] via 10.1.13.3, 00:06:23, Serial0/1
O E2 192.168.5.0/24 [110/65] via 10.1.13.3, 00:06:23, Serial0/1
O E2    10.1.46.0 [110/1] via 10.1.13.3, 00:06:23, Serial0/1
O E2    10.1.45.0 [110/1] via 10.1.13.3, 00:06:23, Serial0/1
O E2 192.168.6.0/24 [110/1] via 10.1.13.3, 00:06:23, Serial0/1
```

On R5

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5#Show ip route ospf | Inc O

O E2 1.0.0.0/8 [110/65] via 10.1.45.4, 00:05:28, Serial0/1
O E2 6.0.0.0/8 [110/1] via 10.1.45.4, 00:05:28, Serial0/1
O E2    10.1.13.0 [110/1] via 10.1.45.4, 00:05:28, Serial0/1
O E2    10.1.46.0 [110/1] via 10.1.45.4, 00:05:28, Serial0/1
O E2 192.168.6.0/24 [110/1] via 10.1.45.4, 00:05:28, Serial0/1
O E2 192.168.1.0/24 [110/65] via 10.1.45.4, 00:05:28, Serial0/1
```

On R6

```
R6#Show ip route rip | Inc R

R    1.0.0.0/8 [120/5] via 10.1.46.4, 00:00:15, FastEthernet0/0
R    5.0.0.0/8 [120/5] via 10.1.46.4, 00:00:15, FastEthernet0/0
R    192.168.5.0/24 [120/5] via 10.1.46.4, 00:00:15, FastEthernet0/0
R       10.1.13.0 [120/5] via 10.1.46.4, 00:00:15, FastEthernet0/0
R       10.1.45.0 [120/5] via 10.1.46.4, 00:00:15, FastEthernet0/0
R    192.168.1.0/24 [120/5] via 10.1.46.4, 00:00:15, FastEthernet0/0
```

Note the routes that are in the same ospf area are showing up as External type 2 routes.

Was this because of the redistribution?
Although redistribution results in this behavior under normal circumstances, this does not follow classic OSPF default behavior. In this case both OSPF enabled routers are configured in the same area (Area 0).

To see the reason behind the behavior:

On R3

```
R3#Show ip bgp vpnv4 vrf CA 1.0.0.0

BGP routing table entry for 1:10:1.0.0.0/8, version 62
Paths: (1 available, best #1, table CA)
  Advertised to update-groups:
        1
  Local
    10.1.13.1 from 0.0.0.0 (3.3.3.3)
      Origin incomplete, metric 65, localpref 100, weight 32768, valid,
sourced, best
      Extended Community: RT:1:105 RT:1:615
        OSPF DOMAIN ID:0x0005:0x000000030200 OSPF RT:0.0.0.0:2:0
        OSPF ROUTER ID:10.1.13.3:0
      mpls labels in/out 306/nolabel
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Every domain-id on a Cisco Router will begin with "0005", "0105" or "0205" which identifies the type of format used. With the "0005" value used the next 4 Octect are viewed as the Global Administrator field and is used as part of the domain-id. The last 2 Octets would be used as a Local Administer field and it's ignored. As you can see by the output, the Global Administrator field is set to "3" which represents the OSPF process-id in HEX. The "0200" output at the end is ignored.
The OSPF RT:0.0.0.0:2:0 value represents the Area of OSPF "0.0.0.0", the internal ospf route "2", and the last value of the last digit "0" means that the route is neither an External 1 nor an External 2 route.

On R4

```
R4#Show ip bgp vpnv4 vrf CA 5.0.0.0

BGP routing table entry for 1:50:5.0.0.0/8, version 65
Paths: (1 available, best #1, table CA)
  Advertised to update-groups:
        1
  Local
    10.1.45.5 from 0.0.0.0 (4.4.4.4)
      Origin incomplete, metric 65, localpref 100, weight 32768, valid,
sourced, best
      Extended Community: RT:1:105 RT:1:615
        OSPF DOMAIN ID:0x0005:0x000000040200 OSPF RT:0.0.0.0:2:0
        OSPF ROUTER ID:10.1.45.4:0
      mpls labels in/out 406/nolabel
```

The domain-id for the prefixes does not match. Under normal OSPF design rules, the process-id of an OSPF router does not need to match any other peer's process-id. However, with MPLS, the router uses the OSPF process as a portion of the domain-id.
Therefore, we must manually change the process-id to match on both sides of the cloud or statically configure a domain-id that matches on both routers.

Routers that do not share the same domain-id are considered type 5 or external to OSPF. If the domain-ids match, the routes are considered type "3" LSAs.

Manually configuring the domain-id is preferred as this does not change the classic understanding of the process-id in OSPF. A clearing of BGP with the "clear ip bgp * in" and "clear ip bgp * out" will speed up the convergence.

On R3

```
R3(config)#router ospf 3 vrf CA
R3(config-router)#Domain-id 0.0.0.1

R3#Clear ip bgp * in
R3#Clear ip bgp * out
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4(config)#router ospf 4 vrf CA
R4(config-router)#Domain-id 0.0.0.1

R4#clear ip bgp * in
R4#clear ip bgp * out
```

On R3

```
R3#Show ip bgp vpnv4 vrf CA 1.0.0.0 | Inc DOMAIN

       OSPF DOMAIN ID:0x0005:0x000000010200 OSPF RT:0.0.0.0:2:0
```

On R4

Domain-ids match

```
R4#Show ip bgp vpnv4 vrf CA 5.0.0.0 | Inc DOMAIN

      OSPF DOMAIN ID:0x0005:0x000000010200 OSPF RT:0.0.0.0:2:0
```

To test and Verify the configuration:

On R1

```
R1#Show ip route ospf | Inc O

O IA 5.0.0.0/8 [110/129] via 10.1.13.3, 00:08:39, Serial0/1
O E2 6.0.0.0/8 [110/1] via 10.1.13.3, 00:39:47, Serial0/1
O IA 192.168.5.0/24 [110/129] via 10.1.13.3, 00:08:39, Serial0/1
O E2    10.1.46.0 [110/1] via 10.1.13.3, 00:39:47, Serial0/1
O IA    10.1.45.0 [110/65] via 10.1.13.3, 00:08:39, Serial0/1
O E2 192.168.6.0/24 [110/1] via 10.1.13.3, 00:39:47, Serial0/1
```

On R5

```
R5#Show ip route ospf | Inc O

O IA 1.0.0.0/8 [110/129] via 10.1.45.4, 00:10:55, Serial0/1
O E2 6.0.0.0/8 [110/1] via 10.1.45.4, 00:40:05, Serial0/1
O IA    10.1.13.0 [110/65] via 10.1.45.4, 00:10:55, Serial0/1
O E2    10.1.46.0 [110/1] via 10.1.45.4, 00:40:05, Serial0/1
O E2 192.168.6.0/24 [110/1] via 10.1.45.4, 00:40:05, Serial0/1
O IA 192.168.1.0/24 [110/129] via 10.1.45.4, 00:10:55, Serial0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Note ONLY the RIP routes are External type 2s and all other OSPF routes are now IA routes.

To verify the configuration;

On R6

```
R6#Show ip route rip

R    1.0.0.0/8 [120/5] via 10.1.46.4, 00:00:10, FastEthernet0/0
R    5.0.0.0/8 [120/5] via 10.1.46.4, 00:00:10, FastEthernet0/0
R    192.168.5.0/24 [120/5] via 10.1.46.4, 00:00:10, FastEthernet0/0
R       10.1.13.0 [120/5] via 10.1.46.4, 00:00:10, FastEthernet0/0
R       10.1.45.0 [120/5] via 10.1.46.4, 00:00:10, FastEthernet0/0
R    192.168.1.0/24 [120/5] via 10.1.46.4, 00:00:10, FastEthernet0/0
```

Since a metric of 5 was set when the routes were redistributed into RIP on R4 (The PE router), all routes end up in R6's routing table with a hop count of 5. However, if R4 was configured with a metric of transparent using "Redistribute BGP 65001 Metric transparent", R6 would have ONLY seen the directly connected routes of R1 and R5, the following tests the metric transparent when redistributing routes into RIPv2.

On R4

```
R4(config)#router rip
R4(config-router)#Address-family ipv4 vrf CB
R4(config-router-af)#NO Redistribute BGP 65001 Metric 5
R4(config-router-af)#Redistribute BGP 65001 Metric transparent
```

On R6

```
R6#Show ip route rip

R       10.1.13.0 [120/1] via 10.1.46.4, 00:00:16, FastEthernet0/0
R       10.1.45.0 [120/1] via 10.1.46.4, 00:00:16, FastEthernet0/0
```

Note ONLY R1 and R5's directly connected links are in R6's routing table, the question is WHY?

On R3

```
R3#Show ip route vrf CA 1.0.0.0

Routing entry for 1.0.0.0/8
  Known via "ospf 3", distance 110, metric 65, type intra area
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Redistributing via bgp 65001
  Advertised by bgp 65001
  Last update from 10.1.13.1 on Serial0/1, 07:35:55 ago
  Routing Descriptor Blocks:
  * 10.1.13.1, from 1.1.1.1, 07:35:55 ago, via Serial0/1
      Route metric is 65, traffic share count is 1
```

On R4

```
R4#Show ip route vrf CA 5.0.0.0

Routing entry for 5.0.0.0/8
  Known via "ospf 4", distance 110, metric 65, type intra area
  Redistributing via bgp 65001
  Advertised by bgp 65001
  Last update from 10.1.45.5 on Serial0/1, 07:22:26 ago
  Routing Descriptor Blocks:
  * 10.1.45.5, from 5.5.5.5, 07:22:26 ago, via Serial0/1
      Route metric is 65, traffic share count is 1
```

Note since the metric of 65 in OSPF is copied into BGP as MED and the MED is copied into RIP as the hop count; from RIP's perspective the routes are 65 hops away, which means that they are inaccessible, the following Debug reveals that R6 receives the routes but it rejects them:

On R6

```
R6#Debug ip rip

RIP: received v2 update from 10.1.46.4 on FastEthernet0/0
 1.0.0.0/8 via 0.0.0.0 in 16 hops  (inaccessible)
 5.0.0.0/8 via 0.0.0.0 in 16 hops  (inaccessible)
 10.1.13.0/24 via 0.0.0.0 in 1 hops
 10.1.45.0/24 via 0.0.0.0 in 1 hops
```

Note the output of the following show command reveals that the directly connected links of R1 and R5 are advertised with a hop count of 1, whereas, the other routes are inaccessible.

```
R4#Show ip route vrf CA 10.1.45.0

Routing entry for 10.1.45.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via bgp 65001
  Advertised by bgp 65001
  Routing Descriptor Blocks:
  * directly connected, via Serial0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
      Route metric is 0, traffic share count is 1

R3#Show ip route vrf CA 10.1.13.0

Routing entry for 10.1.13.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via bgp 65001
  Advertised by bgp 65001
  Routing Descriptor Blocks:
  * directly connected, via Serial0/1
      Route metric is 0, traffic share count is 1
```

On R6

```
R6#Show ip route rip

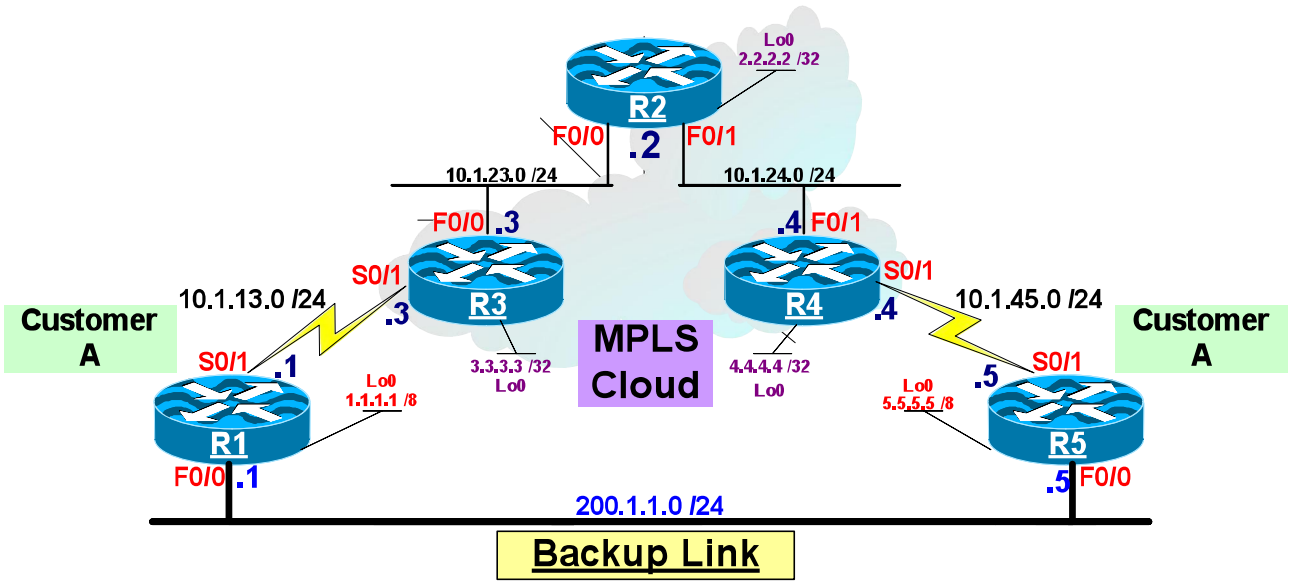
R       10.1.13.0 [120/1] via 10.1.46.4, 00:00:16, FastEthernet0/0
R       10.1.45.0 [120/1] via 10.1.46.4, 00:00:16, FastEthernet0/0
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 3
Backdoor links & OSPF



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 and R5 should be configured in VLAN 100.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the IP addressing according to the above diagram.

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.24.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.24.4 0.0.0.0 area 0
```

To verify the configuration:

On R3

```
R3#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       4.4.4.4 [110/3] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       10.1.24.0 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
```

On R2

```
R2#Show ip route ospf | Inc O

O       3.3.3.3 [110/2] via 10.1.23.3, 00:01:03, FastEthernet0/0
O       4.4.4.4 [110/2] via 10.1.24.4, 00:01:03, FastEthernet0/1
```

On R4

```
R4#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       3.3.3.3 [110/3] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       10.1.23.0 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
```

Task 2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 200 – 299
R3 – 300 – 399
R4 – 400 – 499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 300 399

R3(config)#int f0/0
R3(config-if)#MPLS IP
```

On R2

```
R2(config)#MPLS label protocol LDP

R2(config)#MPLS LDP router-id lo0

R2(config)#MPLS label range 200 299

R2(config)#int F0/0
R2(config-if)#MPLS IP

R2(config-if)#int F0/1
R2(config-if)#MPLS IP
```

On R4

```
R4(config)#MPLS label protocol LDP

R4(config)#MPLS LDP router-id lo0

R4(config)#MPLS label range 400 499

R4(config)#int F0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-if)#MPLS IP
```

To verify the configuration:

On R4

```
R4#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 2.2.2.2.646 - 4.4.4.4.37413
        State: Oper; Msgs sent/rcvd: 9/8; Downstream
        Up time: 00:00:15
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

On R2

```
R2#Show mpls ldp neighbor

    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 3.3.3.3.18975 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 9/10; Downstream
        Up time: 00:01:22
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.3
        Addresses bound to peer LDP Ident:
          10.1.23.3       10.1.13.3       3.3.3.3
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 4.4.4.4.37413 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 8/10; Downstream
        Up time: 00:00:48
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.4
        Addresses bound to peer LDP Ident:
          10.1.24.4       10.1.45.4       4.4.4.4
```

On R3

```
R3#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.18975
        State: Oper; Msgs sent/rcvd: 11/10; Downstream
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        Up time: 00:02:09
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

On R4

```
R4(config)#Router bgp 65001
R4(config-router)#NO BGP default ipv4-unicast
R4(config-router)#Neighbor 3.3.3.3 remote-as 65001
R4(config-router)#Neighbor 3.3.3.3 Update-source Lo0

R4(config-router)#Address-family vpnv4 unicast
R4(config-router-af)#Neighbor 3.3.3.3 Activate
R4(config-router-af)#Neighbor 3.3.3.3 Send-community extended
```

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4 65001      13      13        1    0    0 00:09:57        0
```

Task 4

Configure a **Virtual Routing Forwarding (VRF) Instance with a name of CA** (For Customer A), a **route-distinguisher (RD) of "1:10"** and a **route-target (RT) of "1:100"** on R3. On R4, the same route targets will be used for the vrf, but the **RD should be configured to be "1:20"** and the name of the VRF should be configured as **CB**.

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
R3(config-vrf)#route-target Both 1:100

R3(config)#int S0/1
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip address 10.1.13.3 255.255.255.0
```

On R4

```
R4(config)#ip vrf CB
R4(config-vrf)#RD 1:20
R4(config-vrf)#route-target Both 1:100

R4(config)#int S0/1
R4(config-if)#IP vrf forwarding CB
R4(config-if)#ip address 10.1.45.4 255.255.255.0
```

To verify the VRF Configuration:

On R4

```
R4#Show ip vrf detail

VRF CB; default RD 1:20; default VPNID <not set>
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

On R3

```
R3#Show ip vrf detail

VRF CA; default RD 1:10; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Task 5

Configure OSPF using the same process-id on Customer routers (R1 and R5); configure MP-BGP such that Customer A routers can see each others routes. Customer routers should advertise their Lo0 interface with correct mask.

Let's configure the CE routers:

On R1

```
R1(config)#int lo0
R1(config-if)#ip ospf net point-to-point
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.13.1 0.0.0.0 area 0
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 0

R5(config-router)#int lo0
R5(config-if)#ip ospf net point-to-point
```

Let's configure the PE routers:

On R3

```
R3(config)#Router ospf 1 vrf CA
```

OSPF process 1 already exists and is attached to Default-IP-Routing-Table

Note even though the VRF CA is referenced, the same process ID can NOT be used on the same router.

```
R3(config)#router ospf 2 vrf CA
R3(config-router)#Netw 10.1.13.3 0.0.0.0 area 0
```

You should see the following console message:

%OSPF-5-ADJCHG: Process 2, Nbr 1.1.1.1 on Serial0/1 from LOADING to FULL, Loading Done

To verify the configuration:

On R3

```
R3#Show ip route vrf CA | B Gate
Gateway of last resort is not set

O    1.0.0.0/8 [110/65] via 10.1.13.1, 00:01:49, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
```

On R4

```
R4(config)#Router ospf 2 vrf CB
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 0
```

You should see the following console message:

%OSPF-5-ADJCHG: Process 2, Nbr 5.5.5.5 on Serial0/1 from LOADING to FULL, Loading Done

To verify the configuration:

On R4

```
R4#Show ip route vrf CB | B Gate
Gateway of last resort is not set

O    5.0.0.0/8 [110/65] via 10.1.45.5, 00:00:58, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/1
```

Redistribute OSPF into MP-BGP on PEs:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#Redistribute OSPF 2
```

On R4

```
R4(config)#Router bgp 65001
R4(config-router)#Address-family ipv4 vrf CB
R4(config-router-af)#Redistribute OSPF 2
```

To verify the configuration:

On R3

```
R3#Show ip bgp vpnv4 vrf CA | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                65         32768 ?

*>i5.0.0.0          4.4.4.4                  65    100      0 ?
*> 10.1.13.0/24     0.0.0.0                   0         32768 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4#Show ip bgp vpnv4 vrf CB | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0          3.3.3.3                 65    100      0 ?
*> 5.0.0.0          10.1.45.5               65         32768 ?
*>i10.1.13.0/24     3.3.3.3                  0    100      0 ?
*> 10.1.45.0/24     0.0.0.0                  0         32768 ?
```

Redistribute MP-BGP into OSPF on PEs:

On R3

```
R3(config)#Router ospf 2 vrf CA
R3(config-router)#Redistribute bgp 65001 subnets
```

On R4

```
R4(config)#router ospf 2 vrf CB
R4(config-router)#redistribute BGP 65001 subnets
```

To verify the configuration:

On R1

```
R1#Show ip route ospf | Inc O

O IA 5.0.0.0/8 [110/129] via 10.1.13.3, 00:01:04, Serial0/1
O IA    10.1.45.0 [110/65] via 10.1.13.3, 00:01:04, Serial0/1

R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

On R5

```
R5#Show ip route ospf | Inc O

O IA 1.0.0.0/8 [110/129] via 10.1.45.4, 00:01:20, Serial0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    10.1.13.0 [110/65] via 10.1.45.4, 00:01:20, Serial0/1

R5#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

Task 6

Configure the F0/0 interface of the customer routers (R1 and R5) as a backup link between R1 and R5. Ensure that the MPLS Service is preferred over the backup link. Use the existing IP addressing to accomplish this task.

On R1 and R5

```
Rx(config)#int F0/0
Rx(config-if)#NO shut
```

On R1

```
R1(config)#Router ospf 1
R1(config-router)#Netw 200.1.1.1 0.0.0.0 area 0
```

On R5

```
R5(config)#Router ospf 1
R5(config-router)#Netw 200.1.1.5 0.0.0.0 area 0
```

To verify the routing table:

On R5

```
R5#Show ip route ospf | Inc O

O    1.0.0.0/8 [110/2] via 200.1.1.1, 00:00:01, FastEthernet0/0
O       10.1.13.0 [110/65] via 200.1.1.1, 00:00:01, FastEthernet0/0
```

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip route ospf | Inc O

O     5.0.0.0/8 [110/2] via 200.1.1.5, 00:00:58, FastEthernet0/0
O        10.1.45.0 [110/65] via 200.1.1.5, 00:00:58, FastEthernet0/0
```

Note the routes are reachable via the backup link. Manipulating OSPF cost will NOT change the routing table, because INTRA-area routes are always preferred over INTER-area routes. One way to fix this problem is to configure a Sham link; Sham link was created to resolve this problem. Sham link is an OSPF INTRA-area link configured between the two PE routers. Sham link is included in the SPF calculation, just like any link in OSPF.

When LSAs are flooded across this link all OSPF route types are preserved and they don't have to be converted into INTRA-area or External.

To configure a Sham link:

1. **Create a loopback interface and assign an IP address with a prefix-length of 32.**
2. **Enable VRF forwarding on the loopback interface, in this case "IP vrf forwarding CA" on R3 and "IP vrf forwarding CB" on R4.**
3. **Advertise the /32 IP address of the loopback interface in BGP, on R3 it should be advertised under "Address-family ipv4 vrf CA" and on R4 it should be advertised under "Address-family ipv4 vrf CB".**
4. **Configure Sham link under OSPF, on R3 it should be done under "router ospf 2 vrf CA" and on R4 it should be done under "router ospf 2 vrf CB".**

Step One:

Configure the loopback with a prefix-length of 32:

On R3

```
R3(config-router)#int lo34
R3(config-if)#ip addr 34.1.1.3 255.255.255.255
```

On R4

```
R4(config)#int lo43
R4(config-if)#ip addr 34.1.1.4 255.255.255.255
```

Step Two:

Configure VRF forwarding on the loopback interface:

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#int lo34
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip addr 34.1.1.3 255.255.255.255
```

On R4

```
R4(config)#int lo43
R4(config-if)#ip vrf forwarding CB
R4(config-if)#ip addr 34.1.1.4 255.255.255.255
```

Step Three:

Advertise the ip address of the loopback interface in MP-BGP:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#network 34.1.1.3 mask 255.255.255.255
```

On R4

```
R4(config-if)#router bgp 65001
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#network 34.1.1.4 mask 255.255.255.255
```

Step Four:

Configure the Sham link:

On R3

```
R3(config)#router ospf 2 vrf CA
R3(config-router)#Area 0 Sham-link 34.1.1.3 34.1.1.4 cost 1
```

On R4

```
R4(config)#router ospf 2 vrf CB
R4(config-router)#Area 0 Sham-link 34.1.1.4 34.1.1.3 cost 1
```

You should see the following console message:

Process 2, Nbr 10.1.45.4 on OSPF_SL0 from LOADING to FULL, Loading Done

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3#Show ip ospf sham-links

Sham Link OSPF_SL0 to address 34.1.1.4 is up
Area 0 source address 34.1.1.3
  Run as demand circuit
  DoNotAge LSA allowed. Cost of using 1 State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
    Hello due in 00:00:02
    Adjacency State FULL (Hello suppressed)
    Index 2/2, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

Area 0 was the source for the sham link. This will commonly match the area that has been connected from the PE to the CE. This connection is considered "up" and operational at this time.

As you can see, this connection will run as a demand circuit, The route will remain up indefinately with no interference. The "DoNotAge LSA allowed" should back up this statement with type point-to-point and a cost of 1 as configured under the OSPF process for the VPN. The remaining timers reflect normal OSPF peering.

On R3

```
R3#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           1   FULL/DR         00:00:35    10.1.23.2       FastEthernet0/0
10.1.45.4         0   FULL/  -            -        34.1.1.4        OSPF_SL0
1.1.1.1           0   FULL/  -         00:00:31    10.1.13.1       Serial0/1
```

On R1

```
R1#Show ip route ospf | Inc O

O E2    34.1.1.3 [110/1] via 10.1.13.3, 00:05:50, Serial0/1
O E2    34.1.1.4 [110/1] via 10.1.13.3, 00:05:11, Serial0/1
O    5.0.0.0/8 [110/2] via 200.1.1.5, 00:39:57, FastEthernet0/0
O       10.1.45.0 [110/65] via 200.1.1.5, 00:39:57, FastEthernet0/0
```

On R4

```
R4#Show ip ospf sham-links
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Sham Link OSPF_SL0 to address 34.1.1.3 is up
Area 0 source address 34.1.1.4
  Run as demand circuit
  DoNotAge LSA allowed. Cost of using 1 State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
    Hello due in 00:00:00
    Adjacency State FULL (Hello suppressed)
    Index 2/2, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec

R4#Show ip ospf neighbor

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           1   FULL/DR         00:00:34    10.1.24.2       FastEthernet0/1
10.1.13.3         0   FULL/  -            -        34.1.1.3        OSPF_SL0
5.5.5.5           0   FULL/  -        00:00:38    10.1.45.5       Serial0/1
```

On R5

```
R5#Show ip route ospf | Inc O

O E2     34.1.1.3 [110/1] via 10.1.45.4, 00:08:46, Serial0/1
O E2     34.1.1.4 [110/1] via 10.1.45.4, 00:08:35, Serial0/1
O     1.0.0.0/8 [110/2] via 200.1.1.1, 00:43:09, FastEthernet0/0
O        10.1.13.0 [110/65] via 200.1.1.1, 00:43:09, FastEthernet0/0
```

Note the reason the routers still prefer the backup link is because of the OSPF cost, note the routes are no longer an INTER-area route and they have preserved their true route type.
Since the router sees the routes through the provider and the backup link as INTRA-area routes, the cost can be manipulated such that the routers will take the provider as their primary and the backup link will be used ONLY if the primary link is down.

In order to see the actual cost through the provider, the backup link should be shutdown:

On R1 and R5

```
Rx(config)#Int F0/0
Rx(config-if)#Shut
```

The following show command reveals the OSPF cost through the provider:

On R5

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5#Show ip route ospf | Inc O

O E2     34.1.1.3 [110/1] via 10.1.45.4, 00:00:37, Serial0/1
O E2     34.1.1.4 [110/1] via 10.1.45.4, 00:00:37, Serial0/1
O     1.0.0.0/8 [110/130] via 10.1.45.4, 00:00:37, Serial0/1
O     200.1.1.0/24 [110/130] via 10.1.45.4, 00:00:01, Serial0/1
O        10.1.13.0 [110/129] via 10.1.45.4, 00:00:37, Serial0/1
```

In order to force the provider link as the primary and the FastEthernet link as the backup link, a cost of 131 is assigned to the backup link, as follows:

On Both Routers

```
Rx(config)#Int F0/0
Rx(config-if)#ip ospf cost 131
Rx(config-if)#NO Shut
```

To verify the configuration:

On R1

```
R1#Show ip route ospf | Inc O

O E2     34.1.1.3 [110/1] via 10.1.13.3, 00:00:31, Serial0/1
O E2     34.1.1.4 [110/1] via 10.1.13.3, 00:00:31, Serial0/1
O     5.0.0.0/8 [110/130] via 10.1.13.3, 00:00:31, Serial0/1
O        10.1.45.0 [110/129] via 10.1.13.3, 00:00:31, Serial0/1
```

On R5

```
R5#Show ip route ospf | Inc O

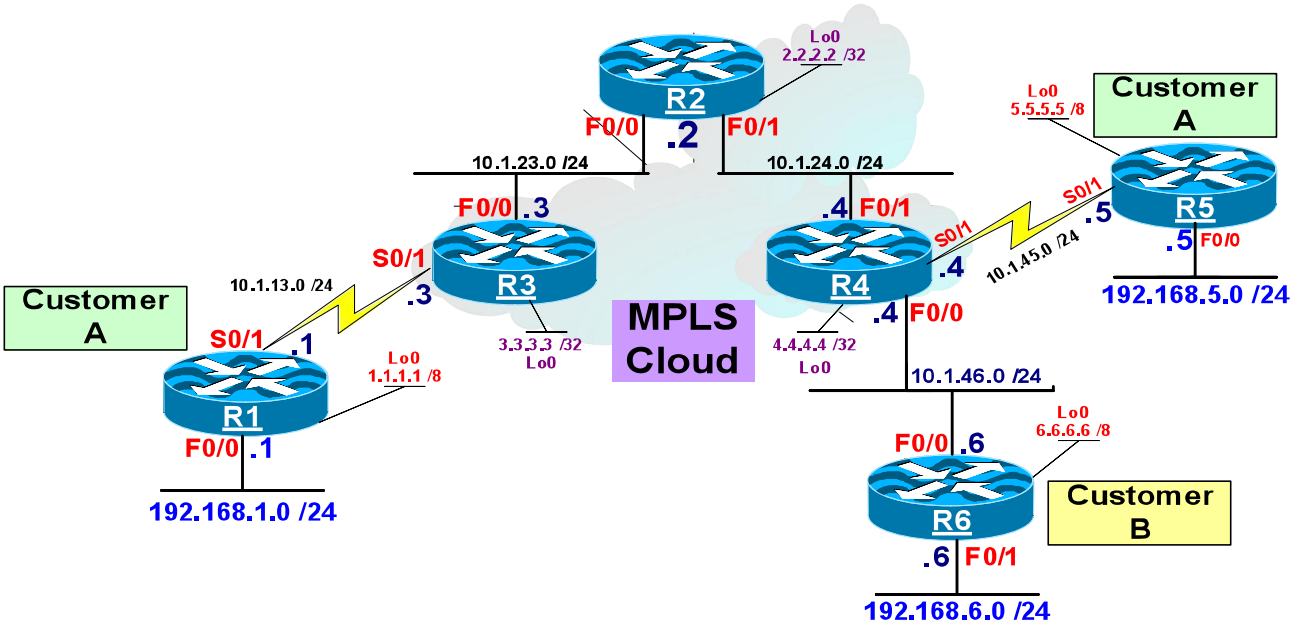
O E2     34.1.1.3 [110/1] via 10.1.45.4, 00:05:24, Serial0/1
O E2     34.1.1.4 [110/1] via 10.1.45.4, 00:05:24, Serial0/1
O     1.0.0.0/8 [110/130] via 10.1.45.4, 00:05:24, Serial0/1
O        10.1.13.0 [110/129] via 10.1.45.4, 00:05:24, Serial0/1
```

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 4
Eigrp Routing in a VPN



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100, R5 in VLAN 500 and F0/1 interface of R6 in VLAN 600.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the F0/1 interface of R6 in VLAN 6.
- Configure the F0/0 interface of R4 and R6 in VLAN 900
- Configure the rest of the routers according to the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.24.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.24.4 0.0.0.0 area 0
```

To verify the configuration:

On R3

```
R3#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       4.4.4.4 [110/3] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       10.1.24.0 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
```

On R2

```
R2#Show ip route ospf | Inc O

O       3.3.3.3 [110/2] via 10.1.23.3, 00:01:03, FastEthernet0/0
O       4.4.4.4 [110/2] via 10.1.24.4, 00:01:03, FastEthernet0/1
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       3.3.3.3 [110/3] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       10.1.23.0 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 200 – 299
R3 – 300 – 399
R4 – 400 – 499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 300 399

R3(config)#int f0/0
R3(config-if)#MPLS IP
```

On R2

```
R2(config)#MPLS label protocol LDP

R2(config)#MPLS LDP router-id lo0

R2(config)#MPLS label range 200 299

R2(config)#int F0/0
R2(config-if)#MPLS IP

R2(config-if)#int F0/1
R2(config-if)#MPLS IP
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#MPLS label protocol LDP

R4(config)#MPLS LDP router-id lo0

R4(config)#MPLS label range 400 499

R4(config)#int F0/1
R4(config-if)#MPLS IP
```

To verify the configuration:

On R4

```
R4#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 2.2.2.2.646 - 4.4.4.4.51493
        State: Oper; Msgs sent/rcvd: 10/8; Downstream
        Up time: 00:00:24
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

On R2

```
R2#Show mpls ldp neighbor

    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 3.3.3.3.20227 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 8/9; Downstream
        Up time: 00:00:40
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.3
        Addresses bound to peer LDP Ident:
          10.1.23.3       10.1.13.3       3.3.3.3
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 4.4.4.4.51493 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 8/10; Downstream
        Up time: 00:00:24
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.4
        Addresses bound to peer LDP Ident:
          10.1.46.4       10.1.24.4       10.1.45.4       4.4.4.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.20227
        State: Oper; Msgs sent/rcvd: 9/8; Downstream
        Up time: 00:00:40
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

On R4

```
R4(config)#Router bgp 65001
R4(config-router)#NO BGP default ipv4-unicast
R4(config-router)#Neighbor 3.3.3.3 remote-as 65001
R4(config-router)#Neighbor 3.3.3.3 Update-source Lo0

R4(config-router)#Address-family vpnv4 unicast
R4(config-router-af)#Neighbor 3.3.3.3 Activate
R4(config-router-af)#Neighbor 3.3.3.3 Send-community extended
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

```
R4#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4 65001      13      13        1    0    0 00:09:57        0
```

Task 4

Configure VRFs on R3 and R4 and enable VRF forwarding on the interfaces of these two routers based on the following chart.

Router	VRF	RD	RT	Interface
R3	CA	1:10	1:105	S0/1
R4	CA	1:50	1:105	S0/1
	CB	1:60	1:600	F0/0

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
R3(config-vrf)#route-target both 1:105

R3(config-vrf)#int S0/1
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip addr 10.1.13.3 255.255.255.0
```

On R4

```
R4(config)#ip vrf CA
R4(config-vrf)#rd 1:50
R4(config-vrf)#route-target both 1:105
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-vrf)#int S0/1
R4(config-if)#ip vrf forwarding CA
R4(config-if)#ip addr 10.1.45.4 255.255.255.0

R4(config)#ip vrf CB
R4(config-vrf)#rd 1:60
R4(config-vrf)#route-target both 1:600

R4(config)#int F0/0
R4(config-if)#ip vrf forwarding CB
R4(config-if)#ip addr 10.1.46.4 255.255.255.0
```

To verify the configuration:

On R3

```
R3#Sh ip vrf det

VRF CA; default RD 1:10; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:105
  Import VPN route-target communities
    RT:1:105
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

On R4

```
R4#Show ip vrf det | S CA

VRF CA; default RD 1:50; default VPNID <not set>
  Interfaces:
    Se0/1

  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:105
  Import VPN route-target communities
    RT:1:105
  No import route-map
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix

R4#Show ip vrf det | S CB

VRF CB; default RD 1:60; default VPNID <not set>
Interfaces:
    Fa0/0
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:600
  Import VPN route-target communities
    RT:1:600
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Task 5

Configure Customer A routers (R1 and R5) with a VRF service that incorporates Eigrp as the routing protocol; Configure the VRF applied to R6 to use OSPF and ensure that this customer participates in the same Closed User Group (CUG) as R1 and R5.
R6 should advertise its loopback0 interface with the correct mask.

To configure the CE Routers:

On R1

```
R1(config)#router eigrp 100
R1(config-router)#No au
R1(config-router)#network 10.1.13.1 0.0.0.0
R1(config-router)#netw 1.1.1.1 0.0.0.0
```

On R5

```
R5(config)#router eigrp 100
R5(config-router)#No au
R5(config-router)#netw 10.1.45.5 0.0.0.0
R5(config-router)#netw 5.5.5.5 0.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R6

```
R6(config-router)#int lo0
R6(config-if)#ip ospf net point-to-point

R6(config)#router ospf 1
R6(config-router)#netw 10.1.46.6 0.0.0.0 are 0
R6(config-router)#netw 6.6.6.6 0.0.0.0 are 0
```

Configuring the PE Routers:

On R3

Configuring Eigrp between PE-CE routers varies from the traditional Eigrp configuration, notice that the definition of Autonomous system number (AS#) is configured under the address-family for the VRF, in this case CA, the AS number configured under the address-family MUST match the AS number configured on the CE router/s. Just like RIPv2 configuration, multiple Eigrp routing context can be configured for every VRF using the address-family command.
One point to remember is that the passive-interface cannot be used in the global instance of EIGRP because it will be inherited to the VPN when the VRF is enabled on the interface and will adversely affect route propagation.

```
R3(config)#Router eigrp 100
R3(config-router)#Address-family ipv4 vrf CA
R3(config-router-af)#No au
R3(config-router-af)#autonomous-system 100
R3(config-router-af)#network 10.1.13.3 0.0.0.0
```



You should get the following console message after adding the "Autonomous-system" command:

%DUAL-5-NBRCHANGE: IP-EIGRP(1) 100: Neighbor 10.1.13.1 (Serial0/1) is up: new adjacency

To verify the configuration:

On R3

```
R3#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

D    1.0.0.0/8 [90/2297856] via 10.1.13.1, 00:05:04, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4(config)#Router eigrp 100
R4(config-router)#address-family ipv4 vrf CA
R4(config-router-af)#No au
R4(config-router-af)#autonomous-system 100
R4(config-router-af)#network 10.1.45.4 0.0.0.0
```

You should see the following console message:

%DUAL-5-NBRCHANGE: IP-EIGRP(1) 100: Neighbor 10.1.45.5 (Serial0/1) is up: new adjacency

To verify the configuration:

On R4

```
R4#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

D     5.0.0.0/8 [90/2297856] via 10.1.45.5, 00:00:59, Serial0/1
      10.0.0.0/24 is subnetted, 1 subnets
C        10.1.45.0 is directly connected, Serial0/1
```

Configuring the PE router (R4) for Customer B (R6):

On R4

```
R4(config)#router ospf 2 vrf CB
R4(config-router)#netw 10.1.46.4 0.0.0.0 area 0
```

To verify the configuration:

On R4

```
R4#Show ip route vrf CB | B Gateway
Gateway of last resort is not set

      6.0.0.0/32 is subnetted, 1 subnets
O        6.6.6.6 [110/2] via 10.1.46.6, 00:00:24, FastEthernet0/0
      10.0.0.0/24 is subnetted, 1 subnets
C        10.1.46.0 is directly connected, FastEthernet0/0
```

Perform a mutual redistribution on the PE and CE routers and verify the routes:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#redistribute Eigrp 100

R3(config)#Router eigrp 100
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#Redistribute BGP 65001 metric 1 1 1 1 1
```

To verify the configuration:

On R3

```
R3#Sh ip bgp vpnv4 vrf CA | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1          2297856         32768 ?
*> 10.1.13.0/24     0.0.0.0                  0         32768 ?
```

On R4

```
R4(config)#Router eigrp 100
R4(config-router)#address-family ipv4 vrf CA
R4(config-router-af)#redistribute BGP 65001 metric 1 1 1 1 1

R4(config)#Router ospf 2 vrf CB
R4(config-router)#Redistribute BGP 65001 subnets

R4(config)#Router bgp 65001
R4(config-router)#address-family ipv4 vrf CA
R4(config-router-af)#Redistribute Eigrp 100
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#Redistribute OSPF 2
```

To verify the configuration:

On R5

```
R5#Sh ip route eigrp | Inc D

D    1.0.0.0/8 [90/2809856] via 10.1.45.4, 00:02:44, Serial0/1
D       10.1.13.0 [90/2681856] via 10.1.45.4, 00:02:44, Serial0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Show ip route eigrp | Inc D

D    5.0.0.0/8 [90/2809856] via 10.1.13.3, 00:01:16, Serial0/1
D       10.1.45.0 [90/2681856] via 10.1.13.3, 00:01:16, Serial0/1
```

On R3

```
R3#Sh ip bgp vpnv4 All | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1          2297856         32768 ?
*>i5.0.0.0          4.4.4.4            2297856    100      0 ?
*> 10.1.13.0/24     0.0.0.0                  0         32768 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
Route Distinguisher: 1:50
*>i5.0.0.0          4.4.4.4            2297856    100      0 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
```

On R4

```
R4#Show ip bgp vpnv4 all | B Net

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10
*>i1.0.0.0          3.3.3.3            2297856    100      0 ?
*>i10.1.13.0/24     3.3.3.3                  0    100      0 ?
Route Distinguisher: 1:50 (default for vrf CA)
*>i1.0.0.0          3.3.3.3            2297856    100      0 ?
*> 5.0.0.0          10.1.45.5          2297856         32768 ?
*>i10.1.13.0/24     3.3.3.3                  0    100      0 ?
*> 10.1.45.0/24     0.0.0.0                  0         32768 ?
Route Distinguisher: 1:60 (default for vrf CB)
*> 6.6.6.6/32       10.1.46.6                2         32768 ?
*> 10.1.46.0/24     0.0.0.0                  0         32768 ?
```

Note R3 contains networks that are associated to VRF CA ONLY, and it does NOT contain any information about VRF CB. Is there something wrong?

When R4 sends all of its routes and their associated route-targets to its PE peer (R3), R3 checks its VRF configurations, and if it does NOT see a given route-target configuration locally, it will simply filter them. It's very similar to VTP pruning. To avoid this behavior, R3 can be configured NOT to filter these routes:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#router bgp 65001
R3(config-router)#NO bgp default route-target filter

R3#Clear ip bgp * in
```

To verify the configuration:

On R3

```
R3#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1           2297856         32768 ?
*>i5.0.0.0          4.4.4.4             2297856    100      0 ?
*> 10.1.13.0/24     0.0.0.0                   0         32768 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
Route Distinguisher: 1:50
*>i5.0.0.0          4.4.4.4             2297856    100      0 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
Route Distinguisher: 1:60
*>i6.6.6.6/32       4.4.4.4                   2    100      0 ?
*>i10.1.46.0/24     4.4.4.4                   0    100      0 ?
```

To test the configuration:

On R1

```
R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/54/56 ms
```

On R5

```
R5#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/56 ms
```

The last configuration in this task requires that R6 to participate in the same Closed User Group (CUG) as R1 and R5:

On R4

```
R4(config)#ip vrf CB
R4(config-vrf)#route-target both 1:105
```

To verify test the configuration:

On R6

```
R6#Sh ip route ospf | Inc O

O E2 1.0.0.0/8 [110/2297856] via 10.1.46.4, 00:00:14, FastEthernet0/0
O E2 5.0.0.0/8 [110/2297856] via 10.1.46.4, 00:00:14, FastEthernet0/0
O E2    10.1.13.0 [110/1] via 10.1.46.4, 00:00:14, FastEthernet0/0
O E2    10.1.45.0 [110/1] via 10.1.46.4, 00:00:14, FastEthernet0/0

R6#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms

R6#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

On R1

```
R1#Show ip route eigrp | Inc D

D    5.0.0.0/8 [90/2809856] via 10.1.13.3, 00:12:07, Serial0/1
D EX    6.6.6.6 [170/2560512256] via 10.1.13.3, 00:01:52, Serial0/1
D EX    10.1.46.0 [170/2560512256] via 10.1.13.3, 00:01:52, Serial0/1
D       10.1.45.0 [90/2681856] via 10.1.13.3, 00:12:07, Serial0/1

R1#Ping 6.6.6.6
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/30/32 ms
```

On R5

```
R5#Show ip route eigrp | Inc D

D    1.0.0.0/8 [90/2809856] via 10.1.45.4, 00:14:08, Serial0/1
D EX    6.6.6.6 [170/2560512256] via 10.1.45.4, 00:02:48, Serial0/1
D       10.1.13.0 [90/2681856] via 10.1.45.4, 00:14:08, Serial0/1
D EX    10.1.46.0 [170/2560512256] via 10.1.45.4, 00:02:48, Serial0/1

R5#Ping 6.6.6.6

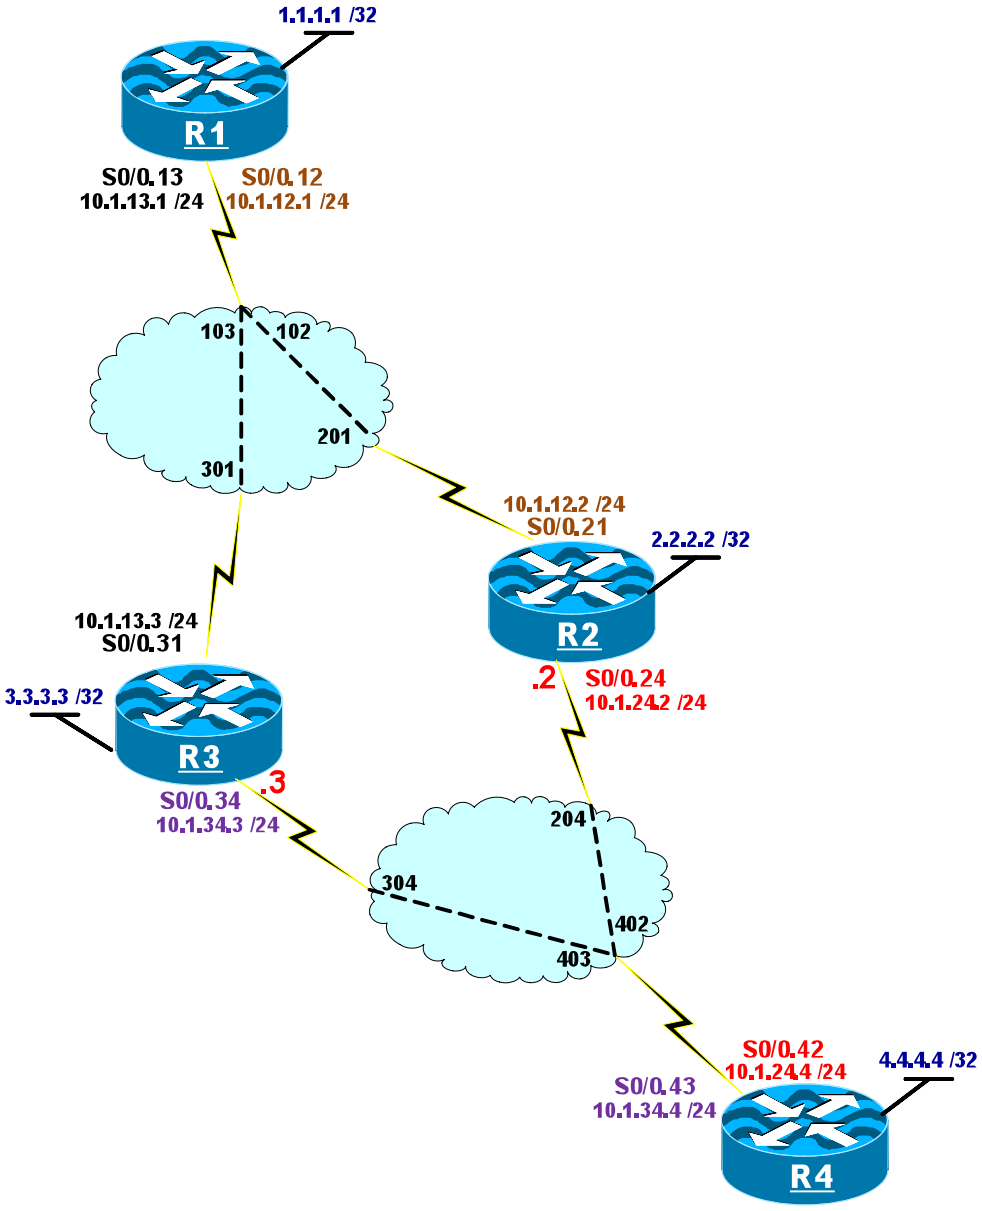
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

LAB 1- EIGRP-SOO



Router	Interface	IP address
R1	Lo0 S0/0.13 S0/0.12	1.1.1.1 /32 10.1.12.1 /24 10.1.13.1/ 24
R2	Lo0 S0/0.21	1.1.1.1/ 32 10.1.21.2 /24

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

	S0/0.24	10.1.24.2 /24
R3	Lo0 S0/0.31 S0/0.34	3.3.3.3/ 32 10.1.13.3/ 24 10.1.34.3/ 24
R4	Lo0 S0/0.42 S0/0.43	4.4.4.4 /32 10.1.24.4/ 24 10.1.34.4/ 24

Lab Setup:

- All Interfaces should be configured in a point-to-point manner.

Task 1

Configure IGP routing between R1 and R2 using ISIS on the loopback and serial interfaces between them. Configure IGP connectivity between R3 and R4 on ALL interfaces.

On R1

```
R1(config)#router isis
R1(config-router)#net 49.0001.1111.0000.1111.00
R1(config-router)#metric-style wide

R1(config-router)#int lo0
R1(config-if)#ip router isis

R1(config-if)#int s0/0.12
R1(config-subif)#ip router isis
```

On R2

```
R2(config)#router isis
R2(config-router)#net 49.0001.2222.0000.2222.00
R2(config-router)#metric-style wide

R2(config-router)#int lo0
R2(config-if)#ip router isis

R2(config-if)#int s0/0.21
R2(config-subif)#ip router isis
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router eigrp 1
R3(config-router)#No auto
R3(config-router)#netw 3.0.0.0
R3(config-router)#netw 10.0.0.0
```

On R4

```
R4(config)#router eigrp 1
R4(config-router)#No auto
R4(config-router)#netw 4.0.0.0
R4(config-router)#netw 10.0.0.0
```

Verifying Connectivity:

R1

```
R1#Sh isis n

System Id      Type Interface   IP Address      State Holdtime Circuit Id
R2             L1L2 Se0/0.12    10.1.12.2       UP    27       00


R1#Ping 2.2.2.2 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/8 ms
```

R3

```
R3#Sh ip eigrp nei

IP-EIGRP neighbors for process 1
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.34.4               Se0/0.34          13 00:02:16    8   300  0  3

R3#Ping 4.4.4.4 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 3.3.3.3
!!!!!
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/12 ms
```

Task 2

Configure Core Routers R1 and R2 to support MPLS VPN's using AS 65012.

On R1

```
R1(config)#router bgp 65012
R1(config-router)#No bgp default ipv4
R1(config-router)#neighbor 2.2.2.2 remot 65012
R1(config-router)#neighbor 2.2.2.2 update lo0

R1(config-router)#address vpnv4
R1(config-router-af)#neighbor 2.2.2.2 act
R1(config-router-af)#neighbor 2.2.2.2 send ext

R1(config-router-af)#interface s0/0.12
R1(config-subif)#mpls ip
```

If LDP is not configured between the Core devices then pings from CE to CE will fail.

On R2

```
R2(config)#router bgp 65012
R2(config-router)#No bgp def ipv4
R2(config-router)#neighbor 1.1.1.1 remot 65012
R2(config-router)#neighbor 1.1.1.1 update lo0

R2(config-router)#address vpnv4
R2(config-router-af)#neighbor 1.1.1.1 act
R2(config-router-af)#neighbor 1.1.1.1 send ext

R2(config-router-af)#interface s0/0.21
R2(config-subif)#mpls ip

R2#Sh ip bgp vpnv4 all summ | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
1.1.1.1         4 65012       3       3        1    0    0 00:00:49        0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Create a VRF called “TST” with unique RD values. Route targets should also be created that match for full connectivity between Customers to be possible. Use EIGRP as the routing intance for PE-CE connectivity.

On R1

```
R1(config)#ip vrf TST
R1(config-vrf)#rd 1.1.1.1:10
R1(config-vrf)#route-target b 65012:100

R1(config-vrf)#int s0/0.13
R1(config-subif)#ip vrf forwarding TST
R1(config-subif)#ip add 10.1.13.1 255.255.255.0

R1(config)#router eigrp 10
R1(config-router)#address-family ipv4 vrf TST
R1(config-router-af)#No auto
R1(config-router-af)#netw 10.0.0.0
R1(config-router-af)#autonomous-system 1
R1(config-router-af)#redistribute bgp 65012 metric 1 1 1 1 1

R1(config-router-af)#router bgp 65012
R1(config-router)#add ipv4 vrf TST
R1(config-router-af)#redistribute eigrp 1
```

On R2

```
R2(config)#ip vrf TST
R2(config-vrf)#rd 2.2.2.2:10
R2(config-vrf)#route-target b 65012:100

R2(config-vrf)#int s0/0.24
R2(config-subif)#ip vrf forwarding TST
R2(config-subif)#ip add 10.1.24.2 255.255.255.0

R2(config)#router eigrp 10
R2(config-router)#address-family ipv4 vrf TST
R2(config-router-af)#No auto
R2(config-router-af)#netw 10.0.0.0
R2(config-router-af)#autonomous-system 1
R2(config-router-af)#redistribute bgp 65012 metric 1 1 1 1 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router-af)#router bgp 65012
R2(config-router)#add ipv4 vrf TST
R2(config-router-af)#redistribute eigrp 1
```

On R3

```
R3#Sh ip route eigrp

     4.0.0.0/32 is subnetted, 1 subnets
D       4.4.4.4 [90/2297856] via 10.1.34.4, 00:00:23, Serial0/0.34
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.24.0 [90/2681856] via 10.1.34.4, 00:00:23, Serial0/0.34
                  [90/2681856] via 10.1.13.1, 00:00:23, Serial0/0.31
```

On R4

```
R4#Sh ip route Eigrp

     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/2297856] via 10.1.34.3, 00:01:07, Serial0/0.43
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.13.0 [90/2681856] via 10.1.34.3, 00:01:07, Serial0/0.43
                  [90/2681856] via 10.1.24.2, 00:01:07, Serial0/0.42
```

Notice how there is duplicate routes for the Serial links that Interconnect the PE and CE routers. The loopback interaces do not show up as duplicates due to the lower cost across the directly connected path.

Task 4

Configure the routers to eliminate the local routes from being learned from the backbone. You cannot create any additional configurations on R3 or R4. No access-list or prefix-list are allowed.

Behavior Prior to filtering Routes:

On R1

```
R1(config)#route-map SOO permit 10
R1(config-route-map)#set extcommunity soo 1:12

R1(config-route-map)#int s0/0.13
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-subif)#ip vrf sitemap SOO
```

Site of Origin (SOO) is an extended community attribute that can be used to prevent routing update loops that can occur as a result of an MPLS VPN design and using a backdoor link.

On R2

```
R2(config)#route-map SOO permit 10
R2(config-route-map)#set extcommunity soo 1:24

R2(config-route-map)#int s0/0.24
R2(config-subif)#ip vrf sitemap SOO
```

Verifying the Configuration:

On R1

```
R1#Sh ip eigrp vrf TST topology 10.1.24.0/24

IP-EIGRP (AS 1): Topology entry for 10.1.24.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2169856
  Routing Descriptor Blocks:
  2.2.2.2, from VPNv4 Sourced, Send flag is 0x0
      Composite metric is (2169856/0), Route is Internal (VPNv4 Sourced)
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 20000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 0
      Extended Community: SoO:1:24
```

The extended Community that was set on R2 affects the traffic that is RECEIVED from R4 and is advertised to R1 as shown above.

On R3

```
R3#Sh ip eigrp top 10.1.24.0/24

IP-EIGRP (AS 1): Topology entry for 10.1.24.0/24
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is 2681856
  Routing Descriptor Blocks:
  10.1.13.1 (Serial0/0.31), from 10.1.13.1, Send flag is 0x0
      Composite metric is (2681856/2169856), Route is Internal
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 40000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
      Extended Community: SoO:1:24

  10.1.34.4 (Serial0/0.34), from 10.1.34.4, Send flag is 0x0
      Composite metric is (2681856/2169856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 40000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
```

R3 also receives this Extended Community in the EIGRP update that was sent by R1. We will discuss why this route was advertised at the end of the section.

```
R3#Sh ip route Eigrp

     4.0.0.0/32 is subnetted, 1 subnets
D       4.4.4.4 [90/2297856] via 10.1.34.4, 00:04:06, Serial0/0.34
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.24.0 [90/2681856] via 10.1.34.4, 00:03:35, Serial0/0.34
                  [90/2681856] via 10.1.13.1, 00:03:35, Serial0/0.31
```

The route shows up in the table due to the redistribution parameters used. It is equated to cost received by R4. R3 will attempt to load balance the traffic due to the default VARIANCE and LOAD SHARE of EIGRP.

On R2

```
R2#Sh ip eigrp vrf TST topology 3.3.3.3/32

IP-EIGRP (AS 1): Topology entry for 3.3.3.3/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2297856
  Routing Descriptor Blocks:
  1.1.1.1, from VPNv4 Sourced, Send flag is 0x0
      Composite metric is (2297856/0), Route is Internal (VPNv4 Sourced)
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 25000 microseconds
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
      Extended Community: SoO:1:12
  10.1.24.4 (Serial0/0.24), from 10.1.24.4, Send flag is 0x0
      Composite metric is (2809856/2297856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 45000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
```

The other Site of Origin value is shown for a complete view of the process. Although this value is RECEIVED by R1 and sent to R2, it is also compared with the local SOO value of 1:24 that has been set. This value does NOT equal the local value; therefore the router will advertise this prefix to R4.

On R4

```
R2#Sh ip eigrp topology 3.3.3.3/32

IP-EIGRP (AS 1): Topology entry for 3.3.3.3/32
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 2297856
  Routing Descriptor Blocks:
  10.1.34.3 (Serial0/0.43), from 10.1.34.3, Send flag is 0x0
      Composite metric is (2297856/128256), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 25000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 1
  10.1.24.2 (Serial0/0.42), from 10.1.24.2, Send flag is 0x0
      Composite metric is (2809856/2297856), Route is Internal
      Vector metric:
        Minimum bandwidth is 1544 Kbit
        Total delay is 45000 microseconds
        Reliability is 255/255
        Load is 1/255
        Minimum MTU is 1500
        Hop count is 2
      Extended Community: SoO:1:12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Sh ip route Eigrp

     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/2297856] via 10.1.34.3, 00:04:47, Serial0/0.43
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.13.0 [90/2681856] via 10.1.34.3, 00:04:11, Serial0/0.43
                  [90/2681856] via 10.1.24.2, 00:04:11, Serial0/0.42
```

The same condition that appeared on R3 is also displayed on R4.

Behavior After filtering Routes:

On R1

```
R1(config)#route-map SOO permit 10
R1(config-route-map)#No set extcommunity soo 1:12
R1(config-route-map)#set extcommunity soo 1:100
```

On R2

```
R2(config)#route-map SOO permit 10
R2(config-route-map)#No set extcommunity soo 1:24
R2(config-route-map)#set extcommunity soo 1:100
```

The SOO community now matches on ALL routes that are being received by both CE routers R3 and R4. R1 and R2 are comparing values that are identical.

On R3

```
R3#Sh ip route Eigrp

     4.0.0.0/32 is subnetted, 1 subnets
D       4.4.4.4 [90/2297856] via 10.1.34.4, 00:01:00, Serial0/0.34
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.24.0 [90/2681856] via 10.1.34.4, 00:00:17, Serial0/0.34
```

On R4

```
R4#Sh ip route Eigrp

     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/2297856] via 10.1.34.3, 00:01:00, Serial0/0.43
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.13.0 [90/2681856] via 10.1.34.3, 00:00:16, Serial0/0.43
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The duplicate routes from the backbone are verified to have been fitered successfully. Remember, this is true for ALL routes.

Task 5

Shut down the interface between R3 and R4 on both routers. Ensure that the Loopback interfaces are reachable from the MPLS backbone for Redundancy purposes ONLY.

On R4

```
R4(config)#int s0/0.43
R4(config-subif)#shut

R4#Sh ip route | b Gateway
Gateway of last resort is not set

     4.0.0.0/32 is subnetted, 1 subnets
C       4.4.4.4 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.24.0 is directly connected, Serial0/0.42
```

Upon shutting down the local interface. The filtering that was configured earlier is still in full effect.

On R3

```
R3(config)#int s0/0.34
R3(config-subif)#shut

R3#Sh ip route | b Gateway
Gateway of last resort is not set

     3.0.0.0/32 is subnetted, 1 subnets
C       3.3.3.3 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.13.0 is directly connected, Serial0/0.31
```

The interface is shutdown on R3 to be consistent.

Option #1:

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#access-list 1 permit 3.3.3.3

R1(config)#route-map SOO permit 5
R1(config-route-map)#match ip address 1
R1(config-route-map)#set extcommunity soo 1:333
```

On R2

```
R2(config)#access-list 1 permit 4.4.4.4
R2(config)#route-map SOO permit 5

R2(config-route-map)#match ip address 1
R2(config-route-map)#set extcommunity soo 1:444
```

The first option is to try and perform the task from the ALLOW perspective. The filter only allows the loopback to be matched and a value of 1:333 is being applied to the 3.3.3.3 prefix. A value of 1:444 is being applied to the 4.4.4.4 prefix. The route-map statement that was added was give a value of "5" to be seen first in the order of statements. The "10" statement is still applied.

Verifying the Configuration:

On R3

```
R3#Sh ip route eigrp

     4.0.0.0/32 is subnetted, 1 subnets
D       4.4.4.4 [90/2809856] via 10.1.13.1, 00:01:32, Serial0/0.31
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.24.0 [90/2681856] via 10.1.13.1, 00:01:32, Serial0/0.31
```

The Result of the configuration is not what you would expect. The loopback prefixe as well as the 10.1.24.0/24 prefix are showing up in the routing table. When the SOO compares values, the order of the statements are indeed viewed by there index. ALL of the prefixes are checked against the first statement and DO NOT match this value. Therefore, both prefixes show up in the routing table.

On R4

```
R4#Sh ip route Eigrp

     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/2809856] via 10.1.24.2, 00:01:57, Serial0/0.42
     10.0.0.0/24 is subnetted, 3 subnets
D       10.1.13.0 [90/2681856] via 10.1.24.2, 00:01:57, Serial0/0.42
```

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#route-map SOO permit 10
R1(config-route-map)#No set extcommunity soo 1:12
R1(config-route-map)#set extcommunity soo 1:100
```

On R2

```
R2(config)#route-map SOO permit 10
R2(config-route-map)#No set extcommunity soo 1:24
R2(config-route-map)#set extcommunity soo 1:100
```

On R3

```
R3#Sh ip route | b Gateway
Gateway of last resort is not set

     3.0.0.0/32 is subnetted, 1 subnets
C       3.3.3.3 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 2 subnets
C       10.1.13.0 is directly connected, Serial0/0.31
C       10.1.34.0 is directly connected, Serial0/0.34
```

The original configured is restored to implement the true solution to the task.

Option #2:

On R1

```
R1(config)#No route-map SOO permit 5

R1(config)#access-list 2 permit 10.1.13.0 0.0.0.255

R1(config)#route-map SOO permit 10
R1(config-route-map)#match ip address 2
```

The problematic Route-map stement is removed and the other statement is changed to add the SOO 1:100 vlaue to the Seria link only. This is a NOT-ALLOW perspective.

On R2

```
R2(config)#No route-map SOO permit 5

R2(config)#access-list 2 permit 10.1.24.0 0.0.0.255

R2(config)#route-map SOO permit 10
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-route-map)#match ip address 2
```

On R3

```
R3#Sh ip route Eigrp

     4.0.0.0/32 is subnetted, 1 subnets
D       4.4.4.4 [90/2809856] via 10.1.13.1, 00:00:04, Serial0/0.31
```

On R4

```
R4#Sh ip route Eigrp

     3.0.0.0/32 is subnetted, 1 subnets
D       3.3.3.3 [90/2809856] via 10.1.24.2, 00:00:22, Serial0/0.42

R4#Ping 3.3.3.3 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
Packet sent with a source address of 4.4.4.4
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/3/8 ms
```

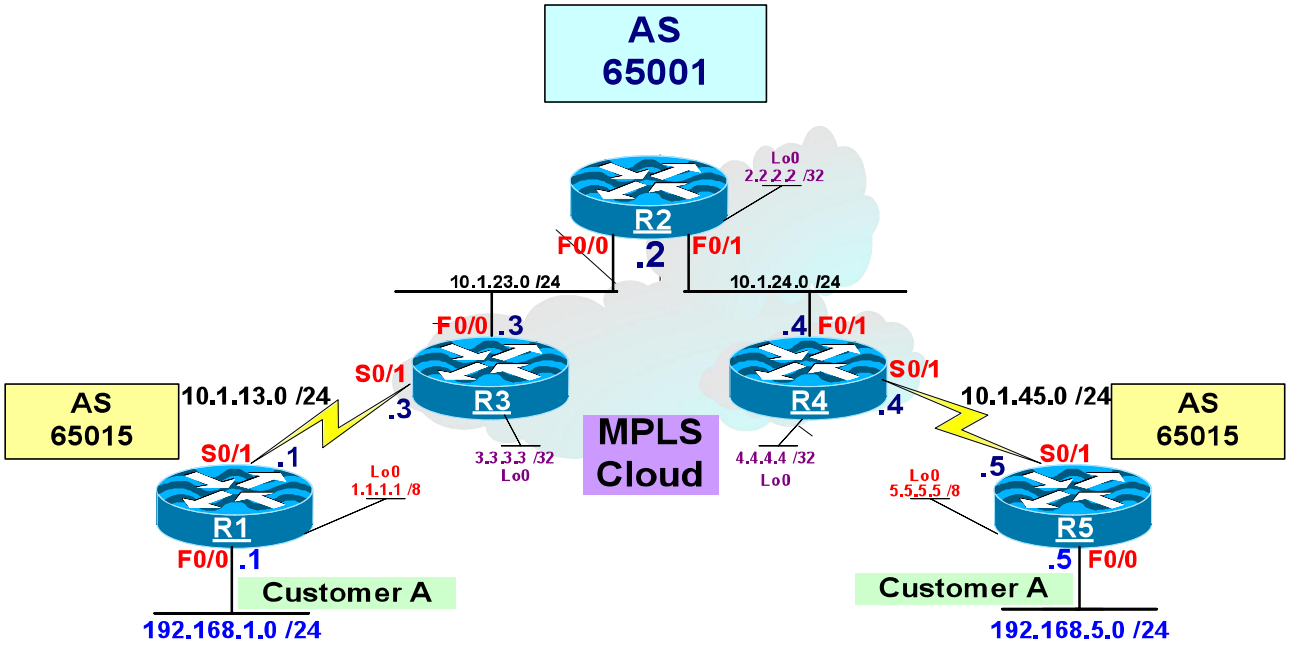
The desired results are achieved and the ping verifies the connectivity requirement.

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 6
BGP Routing in a VPN



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100 and R5 in VLAN 500.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the rest of the routers according to the above diagram.

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.24.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.24.4 0.0.0.0 area 0
```

To verify the configuration:

On R3

```
R3#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       4.4.4.4 [110/3] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       10.1.24.0 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
```

On R2

```
R2#Show ip route ospf | Inc O

O       3.3.3.3 [110/2] via 10.1.23.3, 00:01:03, FastEthernet0/0
O       4.4.4.4 [110/2] via 10.1.24.4, 00:01:03, FastEthernet0/1
```

On R4

```
R4#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       3.3.3.3 [110/3] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       10.1.23.0 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 200 – 299
R3 – 300 – 399
R4 – 400 – 499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 300 399

R3(config)#int f0/0
R3(config-if)#MPLS IP
```

On R2

```
R2(config)#MPLS label protocol LDP

R2(config)#MPLS LDP router-id lo0

R2(config)#MPLS label range 200 299

R2(config)#int F0/0
R2(config-if)#MPLS IP

R2(config-if)#int F0/1
R2(config-if)#MPLS IP
```

On R4

```
R4(config)#MPLS label protocol LDP

R4(config)#MPLS LDP router-id lo0

R4(config)#MPLS label range 400 499

R4(config)#int F0/1
R4(config-if)#MPLS IP
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R4

```
R4#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 2.2.2.2.646 - 4.4.4.4.57906
        State: Oper; Msgs sent/rcvd: 9/8; Downstream
        Up time: 00:00:05
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

On R2

```
R2#Show mpls ldp neighbor

    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 3.3.3.3.41759 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 8/9; Downstream
        Up time: 00:00:22
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.3
        Addresses bound to peer LDP Ident:
          10.1.23.3       10.1.13.3       3.3.3.3
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 4.4.4.4.57906 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 8/9; Downstream
        Up time: 00:00:05
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.4
        Addresses bound to peer LDP Ident:
          10.1.24.4       10.1.45.4       4.4.4.4
```

On R3

```
R3#Show mpls ldp neighbor

        Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.41759
        State: Oper; Msgs sent/rcvd: 9/8; Downstream
        Up time: 00:00:22
        LDP discovery sources:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    FastEthernet0/0, Src IP addr: 10.1.23.2
  Addresses bound to peer LDP Ident:
    10.1.23.2       10.1.24.2       2.2.2.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

On R4

```
R4(config)#Router bgp 65001
R4(config-router)#NO BGP default ipv4-unicast
R4(config-router)#Neighbor 3.3.3.3 remote-as 65001
R4(config-router)#Neighbor 3.3.3.3 Update-source Lo0

R4(config-router)#Address-family vpnv4 unicast
R4(config-router-af)#Neighbor 3.3.3.3 Activate
R4(config-router-af)#Neighbor 3.3.3.3 Send-community extended
```

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4 65001      13      13        1    0    0 00:09:57        0
```

Task 4

Configure a **Virtual Routing Forwarding (VRF) Instance with a name of CA** (For Customer A), a **route-distinguisher (RD) of "1:10"** and a **route-target (RT) of "1:100"** on R3. On R4, the same route targets will be used for the vrf, but the **RD should be configured to be "1:20"** and the name of the VRF should be configured as **CB**.

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#RD 1:10
R3(config-vrf)#route-target Both 1:100

R3(config)#int S0/1
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip address 10.1.13.3 255.255.255.0
```

On R4

```
R4(config)#ip vrf CB
R4(config-vrf)#RD 1:20
R4(config-vrf)#route-target Both 1:100

R4(config)#int S0/1
R4(config-if)#IP vrf forwarding CB
R4(config-if)#ip address 10.1.45.4 255.255.255.0
```

To verify the VRF Configuration:

On R4

```
R4#Show ip vrf detail

VRF CB; default RD 1:20; default VPNID <not set>

  Interfaces:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

On R3

```
R3#Show ip vrf detail

VRF CA; default RD 1:10; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Task 5

Configure BGP as the MPLS Routing context between the CE and their respective PE (R1 and R3 as well as R4 and R5). The customer AS of 65015 should be assigned to both customer sites, whereas, the provider should use AS 65001.

Configuring the CE routers:

On R1

```
R1(config)#router bgp 65015
R1(config-router)#No au
R1(config-router)#neigh 10.1.13.3 remot 65001
R1(config-router)#Netw 1.0.0.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-router)#Netw 192.168.1.0
```

On R5

```
R5(config)#router bgp 65015
R5(config-router)#No au
R5(config-router)#Neigh 10.1.45.4 remot 65001
R5(config-router)#Netw 5.0.0.0
R5(config-router)#Netw 192.168.5.0
```

Configuring the PE routers to establish a peer session with the CE routers:

Note, no redistribution is necessary for MP-BGP. The peering is identical to that of Normal BGP except it is in the VPN instead. However, normal BGP rules apply within the VPN.

On R4

```
R4(config)#router bgp 65001
R4(config-router)#No au
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#neighbor 10.1.45.5 remot 65015
```

You should see the following console messages before moving on:

```
%BGP-5-ADJCHANGE: neighbor 10.1.45.5 vpn vrf CB Up
```

On R3

```
R3(config)#router bgp 65001
R3(config-router)#No au
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#Neighbor 10.1.13.1 remot 65015
```

You should see the following console messages before moving on:

```
%BGP-5-ADJCHANGE: neighbor 10.1.13.1 vpn vrf CA Up
```

To verify the configuration

On R5

```
R5#Sh ip bgp summ | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.45.4       4 65001      36      36        3    0    0 00:31:48        0
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

On R1

```
R1#Show ip bgp summ | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.13.3       4 65001      51      51        3    0    0 00:46:31        0
```

On R3

```
R3#Sh ip bgp vpnv4 vrf CA | B Net

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                0             0 65015 i
*>i5.0.0.0          4.4.4.4                  0    100      0 65015 i
*> 192.168.1.0      10.1.13.1                0             0 65015 i
*>i192.168.5.0      4.4.4.4                  0    100      0 65015 i
```

On R4

```
R4#Sh ip bgp vpnv4 vrf CB | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0          3.3.3.3                  0    100      0 65015 i
*> 5.0.0.0          10.1.45.5                0             0 65015 i
*>i192.168.1.0      3.3.3.3                  0    100      0 65015 i
*> 192.168.5.0      10.1.45.5                0             0 65015 i
```

On R1

```
R1#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 192.168.1.0      0.0.0.0                  0         32768 i
```

On R5

```
R5#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path
*> 5.0.0.0          0.0.0.0                  0         32768 i
*> 192.168.5.0      0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Although the routes were verified to have been received via MP-BGP and redistribution is not a problem that can exist here, the routes are not learned by the CE routers. This is a result of a normal BGP process of loop prevention. There are two ways to resolve this matter:

1. **Configuring the CE solution**
2. **Configuring the PE solution**

Configuring the CE routers to resolve the problem:

On R1

```
R1(config)#router bgp 65015
R1(config-router)#neighbor 10.1.13.3 allowas-in
```

This feature can be used as an alternative if the Service Provider controls the CE router/s.
When network 5.0.0.0 /8 is advertised by the remote office (R5) the AS number is prepended (65015), which is used as a loop prevention mechanism in BGP. The prefix is received by the provider and then, advertised to R1, which is in AS 65015 (Same as R5's AS), when R1 receives the route, it sees its own AS number in the AS-Path attribute and it discards the route as a loop avoidance mechanism, the "allowas-in" command allows the local router to receive routes even if the local router sees its own AS in the AS-Path attribute.

To verify the configuration:

On R1

Note R5 ignored its own AS number in the inbound updates

```
R1#Show ip bgp | B Network
   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          10.1.13.3                              0 65001 65015 i
*> 192.168.1.0      0.0.0.0                  0         32768 i
*> 192.168.5.0      10.1.13.3                              0 65001 65015 i
```

The same feature must be configured on R5:

On R5

```
R5(config)#router bgp 65015
R5(config-router)#Neighbor 10.1.45.4 allowas-in
```

To verify the configuration:

On R5

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 65001 65015 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
*> 192.168.1.0      10.1.45.4                              0 65001 65015 i
*> 192.168.5.0      0.0.0.0                  0         32768 i
```



To configure the PE routers to resolve the problem:

On R5

```
R5(config)#router bgp 65015
R5(config-router)#NO Neighbor 10.1.45.4 allowas-in
```

On R1

```
R1(config)#router bgp 65015
R1(config-router)#NO neighbor 10.1.13.3 allowas-in
```

On R3

```
R3(config)#router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#Neighbor 10.1.13.1 as-override
```

You should receive the following console messages:

%BGP-5-ADJCHANGE: neighbor 10.1.13.1 vpn vrf CA Down AS-override change
%BGP-5-ADJCHANGE: neighbor 10.1.45.5 vpn vrf CA Up

On R4

```
R4(config)#router bgp 65001
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#Neighbor 10.1.45.5 as-override
```

You should receive the following console messages:

%BGP-5-ADJCHANGE: neighbor 10.1.45.5 vpn vrf CB Down AS-override change
%BGP-5-ADJCHANGE: neighbor 10.1.45.5 vpn vrf CB Up

To verify the configuration:

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Note the AS number of the provided is prepended

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 5.0.0.0          10.1.13.3                              0 65001 65001 i
*> 192.168.1.0      0.0.0.0                  0         32768 i
*> 192.168.5.0      10.1.13.3                              0 65001 65001 i
```

On R5

```
R5#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 65001 65001 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
*> 192.168.1.0      10.1.45.4                              0 65001 65001 i
*> 192.168.5.0      0.0.0.0                  0         32768 i
```

Unlike the previous solution where the AS number was ignored, using this feature, the provider replaces the customer's AS number with its own. The customer MUST allow prepended AS numbers or else they will NOT get the BGP routes. Because from the customers perspective, it looks like the provider prepended its own AS number.

Task 6

Configure a backup link on R5 to R4 using frame-relay point-to-point configuration, you should use S0/0.45 with an IP address of 10.1.100.4 /24 on the PE router (R4) and S0/0.54 with an IP address of 10.1.100.5 /24 on the CE router (R5), this configuration will provide a dual Home with the provider, configure this task such that the new frame-relay Point-to-point link is used as the primary link outbound.

Step One:

Configuring the Frame-relay connection between R4 and R5:

On R5

```
R5(config)#Int S0/0
R5(config-if)#encap frame

R5(config-if)#int S0/0.54 p
R5(config-subif)#ip addr 10.1.100.5 255.255.255.0
R5(config-subif)#frame interface-dlci 504
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-fr-dlci)#int S0/0
R5(config-if)#No shut
```

On R4

```
R4(config)#int S0/0
R4(config-if)#encap frame

R4(config-if)#int s0/0.45 P
R4(config-subif)#ip addr 10.1.100.4 255.255.255.0
R4(config-subif)#frame interface-dlci 405

R4(config-fr-dlci)#int s0/0
R4(config-if)#No shut
```

To verify the configuration:

On R4

```
R4#Ping 10.1.100.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.100.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/61 ms
```

Step Two:

Configure VRF forwarding on the Frame-relay interfaces of R4 to R5:

On R4

```
R4(config)#int S0/0.45

R4(config-subif)#ip vrf forwarding CB
R4(config-subif)#ip addr 10.1.100.4 255.255.255.0
```

Step Three:

Establish a BGP peer session between R4 and R5 using the new link:

On R5

```
R5(config)#router bgp 65015
R5(config-router)#Neighbor 10.1.100.4 remote-as 65001
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4(config)#router bgp 65001
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#Neighbor 10.1.100.5 remote-as 65015
```

You should see the following console message before moving on:

```
%BGP-5-ADJCHANGE: neighbor 10.1.100.5 vpn vrf CB Up
```

To verify the configuration:

On R5

```
R5#Sh ip bgp summary | b Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.45.4       4 65001      90      88        9    0    0 00:59:28        2
10.1.100.4      4 65001       7       7        9    0    0 00:01:14        0
```

Note the local router does NOT maintain any routes from its neighboring BGP router 10.1.100.4, because the loop avoidance mechanism is NOT allowing the local router to have any route/s from its remote location router (R1), to fix this problem, the "allowas-in" command is used on the CE router (R5):

On R5

```
R5(config)#router bgp 65015
R5(config-router)#Neighbor 10.1.100.4 allowas-in
```

To verify the configuration:

On R5

```
R5#Sh ip bgp summary | b Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.45.4       4 65001      92      90        9    0    0 01:01:53        2
10.1.100.4      4 65001      11      10        9    0    0 00:03:39        4

R5#Show ip bgp | B Network

     Network          Next Hop            Metric LocPrf Weight Path
*    1.0.0.0          10.1.100.4                             0 65001 65015 i
*>                    10.1.45.4                              0 65001 65001 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*  5.0.0.0          10.1.100.4                             0 65001 65015 i
*>                  0.0.0.0                  0         32768 i
*  192.168.1.0      10.1.100.4                             0 65001 65015 i
*>                  10.1.45.4                              0 65001 65001 i
*  192.168.5.0      10.1.100.4                             0 65001 65015 i
*>                  0.0.0.0                  0         32768 i
```

The output of above shows command reveals that network "5.0.0.0" that resides on R5 is being learned from the core. This network was advertised to the provider by the local router through its F0/0 interface, the provider then advertised it back to the local router through the Frame-relay interface, because the new peer session is configured with the "Allowas-in" command, it was relearned through the provider.

Therefore the routes that are local to "65015" are being re-announced back into the AS as a valid route path. Although BGP route selection favors routes that are locally generated over routes that are external, this route should not have entered in the BGP table at all.

The solution to this problem is configuration of "SOO" which stands for Site Of Origin, as follows:

Step One:

Configure a route-map and set an extended community of soo 1:105

On R4

```
R4(config)#route-map TST permit 10
R4(config-route-map)#set extcommunity soo 1:105
```

Step Two:

Assign the route-map from the previous step to the neighbors inbound

```
R4(config)#router bgp 65001
R4(config-router)#address-family ipv4 vrf CB

R4(config-router-af)#Neighbor 10.1.45.5 route-map TST in
R4(config-router-af)#Neighbor 10.1.100.5 route-map TST in
```

You must "Clear ip BGP * in" for this policy to take effect:

```
R4#Clear ip bgp * in
```

Note the following reveals that the extended community of "soo 1:105" is attached to network 5.0.0.0 through both neighbors:

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4#Show ip bgp vpnv4 vrf CB 5.0.0.0

BGP routing table entry for 1:20:5.0.0.0/8, version 21
Paths: (2 available, best #2, table CB)
Flag: 0xA00
  Advertised to update-groups:
        1
  65015
    10.1.100.5 from 10.1.100.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, external
      Extended Community: SoO:1:105 RT:1:100
      mpls labels in/out 404/nolabel
  65015
    10.1.45.5 from 10.1.45.5 (5.5.5.5)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: SoO:1:105 RT:1:100
      mpls labels in/out 404/nolabel
```

Step Three:

An extended community-list is configured to identify the extended community of "soo 1:105":

```
R4(config)#IP extcommunity-list 1 permit soo 1:105
```

Step Four:

A route-map called "TEST is configured to deny the extended community-list configured in the previous step and permit everything else

```
R4(config)#Route-map TEST deny 10
R4(config-route-map)#match extcommunity 1

R4(config)#Route-map TEST permit 90
```

Step Five:

The route-map is assigned to the two neighbors outbound:

```
R4(config)#router bgp 65001
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#Neighbor 10.1.100.5 route-map TEST out
R4(config-router-af)#Neighbor 10.1.45.5 route-map TEST out

R4#Clear ip bgp * Out
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R5

```
R5#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0          10.1.100.4                             0 65001 65015 i
*>                  10.1.45.4                              0 65001 65001 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
*  192.168.1.0      10.1.100.4                             0 65001 65015 i
*>                  10.1.45.4                              0 65001 65001 i
*> 192.168.5.0      0.0.0.0                  0         32768 i
```

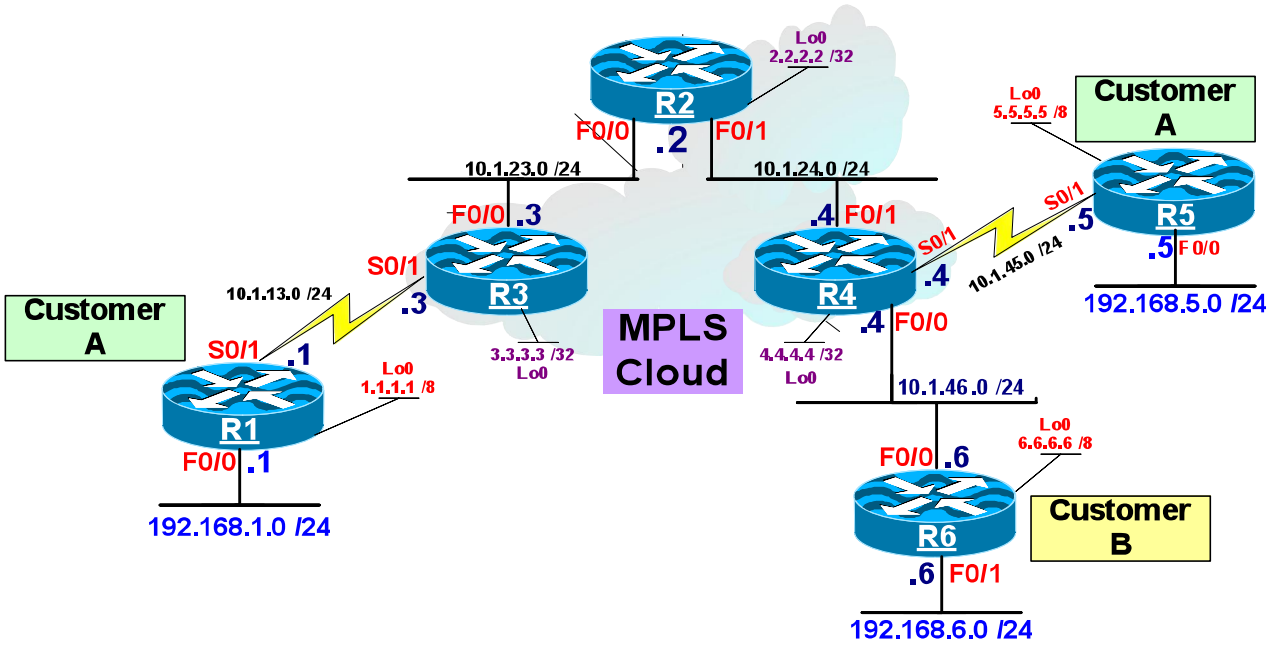
Note networks that were originated by R5 is no longer being learned through the provider.

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 7
Complex VPN'S and Filters



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100, R5 in VLAN 500 and F0/1 interface of R6 in VLAN 600.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the F0/0 interface of R4 and R6 in VLAN 900
- Configure the rest of the routers according to the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 10.1.23.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.24.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.24.4 0.0.0.0 area 0
```

To verify the configuration:

On R3

```
R3#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       4.4.4.4 [110/3] via 10.1.23.2, 00:00:24, FastEthernet0/0
O       10.1.24.0 [110/2] via 10.1.23.2, 00:00:24, FastEthernet0/0
```

On R2

```
R2#Show ip route ospf | Inc O

O       3.3.3.3 [110/2] via 10.1.23.3, 00:01:03, FastEthernet0/0
O       4.4.4.4 [110/2] via 10.1.24.4, 00:01:03, FastEthernet0/1
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Show ip route ospf | Inc O

O       2.2.2.2 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       3.3.3.3 [110/3] via 10.1.24.2, 00:01:41, FastEthernet0/1
O       10.1.23.0 [110/2] via 10.1.24.2, 00:01:41, FastEthernet0/1
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 200 – 299
R3 – 300 – 399
R4 – 400 – 499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 300 399

R3(config)#int f0/0
R3(config-if)#MPLS IP
```

On R2

```
R2(config)#MPLS label protocol LDP

R2(config)#MPLS LDP router-id lo0

R2(config)#MPLS label range 200 299

R2(config)#int F0/0
R2(config-if)#MPLS IP

R2(config-if)#int F0/1
R2(config-if)#MPLS IP
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#MPLS label protocol LDP

R4(config)#MPLS LDP router-id lo0

R4(config)#MPLS label range 400 499

R4(config)#int F0/1
R4(config-if)#MPLS IP
```

To verify the configuration:

On R4

```
R4#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 4.4.4.4:0
        TCP connection: 2.2.2.2.646 - 4.4.4.4.64114
        State: Oper; Msgs sent/rcvd: 70/65; Downstream
        Up time: 00:51:15
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

On R2

```
R2#Show mpls ldp neighbor

    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 3.3.3.3.34077 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 67/68; Downstream
        Up time: 00:51:45
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.3
        Addresses bound to peer LDP Ident:
          10.1.23.3       10.1.13.3       3.3.3.3
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 4.4.4.4.64114 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 65/70; Downstream
        Up time: 00:51:15
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 10.1.24.4
        Addresses bound to peer LDP Ident:
          10.1.46.4       10.1.24.4       10.1.45.4       4.4.4.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.34077
        State: Oper; Msgs sent/rcvd: 68/67; Downstream
        Up time: 00:51:45
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

On R4

```
R4(config)#Router bgp 65001
R4(config-router)#NO BGP default ipv4-unicast
R4(config-router)#Neighbor 3.3.3.3 remote-as 65001
R4(config-router)#Neighbor 3.3.3.3 Update-source Lo0

R4(config-router)#Address-family vpnv4 unicast
R4(config-router-af)#Neighbor 3.3.3.3 Activate
R4(config-router-af)#Neighbor 3.3.3.3 Send-community extended
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

```
R4#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4 65001      13      13        1    0    0 00:09:57        0
```

Task 4

Configure VRFs on R3 and R4 and enable VRF forwarding on the interfaces of these two routers based on the following chart.

Router	VRF	RD	RT	Interface
R3	CA	1:10	1:105	S0/1
R4	CA	1:50	1:105	S0/1
	CB	1:60	1:600	F0/0

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
R3(config-vrf)#route-target both 1:105

R3(config-vrf)#int S0/1
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip addr 10.1.13.3 255.255.255.0
```

On R4

```
R4(config)#ip vrf CA
R4(config-vrf)#rd 1:50
R4(config-vrf)#route-target both 1:105
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-vrf)#int S0/1
R4(config-if)#ip vrf forwarding CA
R4(config-if)#ip addr 10.1.45.4 255.255.255.0

R4(config)#ip vrf CB
R4(config-vrf)#rd 1:60
R4(config-vrf)#route-target both 1:600

R4(config)#int F0/0
R4(config-if)#ip vrf forwarding CB
R4(config-if)#ip addr 10.1.46.4 255.255.255.0
```

To verify the configuration:

On R3

```
R3#Sh ip vrf det

VRF CA; default RD 1:10; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:105
  Import VPN route-target communities
    RT:1:105
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

On R4

```
R4#Show ip vrf det | S CA

VRF CA; default RD 1:50; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:105
  Import VPN route-target communities
    RT:1:105
  No import route-map
  No export route-map
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix

R4#Show ip vrf det | S CB

VRF CB; default RD 1:60; default VPNID <not set>
Interfaces:
    Fa0/0
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:1:600
  Import VPN route-target communities
    RT:1:600
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Task 5

Customer A on R1 should run OSPF using process id of "1" as the IGP, this router should advertise its loopback 0 interface with the correct mask; Customer A on R5 should run EIGRP with an AS number of "100", and R6 should run BGP with an AS number of 65006. Configure MP-BGP accordingly; these routers should participate in the same Closed User Group (CUG).

To configure R1 and its PE (R3):

On R1

```
R1(config)#router ospf 1
R1(config-router)#Network 10.1.13.1 0.0.0.0 area 0
R1(config-router)#Network 1.1.1.1 0.0.0.0 area 0

R1(config)#int lo0
R1(config-if)#ip ospf Network point-to-point
```

On R3

```
R3(config)#router ospf 2 vrf CA
R3(config-router)#Network 10.1.13.3 0.0.0.0 area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

You should see the following console message:

```
%OSPF-5-ADJCHG: Process 2, Nbr 1.1.1.1 on Serial0/1 from LOADING to FULL,
Loading Done
```

To verify the configuration:

On R3

```
R3#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

O    1.0.0.0/8 [110/65] via 10.1.13.1, 00:01:25, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
```

To configure R5 and its PE (R4):

On R5

```
R5(config)#router eigrp 100
R5(config-router)#No au
R5(config-router)#Netw 10.1.45.5 0.0.0.0
R5(config-router)#Netw 5.5.5.5 0.0.0.0
```

On R4

```
R4(config)#router eigrp 100
R4(config-router)#No au
R4(config-router)#address-family ipv4 vrf CA
R4(config-router-af)#autonomous-system 100
R4(config-router-af)#No au
R4(config-router-af)#Netw 10.1.45.4 0.0.0.0
```

You should see the following console message:

%DUAL-5-NBRCHANGE: IP-EIGRP(1) 100: Neighbor 10.1.45.5 (Serial0/1) is up: new adjacency

To verify the configuration:

On R4

```
R4#Show ip route vrf CA | B Gateway
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
Gateway of last resort is not set

D     5.0.0.0/8 [90/2297856] via 10.1.45.5, 00:00:57, Serial0/1
      10.0.0.0/24 is subnetted, 1 subnets
C        10.1.45.0 is directly connected, Serial0/1
```

To configure R6 and its PE (R4):

On R6

```
R6(config)#router bgp 65006
R6(config-router)#neighbor 10.1.46.4 remote-as 65001
R6(config-router)#Netw 6.0.0.0
```

On R4

```
R4(config)#router bgp 65001
R4(config-router)#address-family ipv4 vrf CB
R4(config-router-af)#Neighbor 10.1.46.6 remote-as 65006
```

To verify the configuration:

On R4

```
R4#Show ip bgp vpnv4 vrf CB | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:60 (default for vrf CB)
*> 6.0.0.0          10.1.46.6                0             0 65006 i
```

In the following step OSPF is redistributed into BGP:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#Redistribute ospf 2
```

To verify the configuration:

On R3

```
R3#Sh ip bgp vpnv4 all | B Network
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                65         32768 ?
*> 10.1.13.0/24     0.0.0.0                   0         32768 ?
```

In the following step Eigrp is redistributed into the BGP:

On R4

```
R4(config)#router bgp 65001
R4(config-router)#address-family ipv4 vrf CA
R4(config-router-af)#Redistribute Eigrp 100
```

To verify the configuration:

On R4

```
R4#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10
*>i1.0.0.0          3.3.3.3                  65    100      0 ?
*>i10.1.13.0/24     3.3.3.3                   0    100      0 ?
Route Distinguisher: 1:50 (default for vrf CA)
*>i1.0.0.0          3.3.3.3                  65    100      0 ?
*> 5.0.0.0          10.1.45.5           2297856         32768 ?
*>i10.1.13.0/24     3.3.3.3                   0    100      0 ?
*> 10.1.45.0/24     0.0.0.0                   0         32768 ?
Route Distinguisher: 1:60 (default for vrf CB)
*> 6.0.0.0          10.1.46.6                 0             0 65006 i
```

In the following step BGP is redistributed into the Customer IGP:

On R3

```
R3(config)#router ospf 2 vrf CA
R3(config-router)#redistribute BGP 65001 subnets
```

On R4

```
R4(config)#router eigrp 100
R4(config-router)#address-family ipv4 vrf CA
R4(config-router-af)#Redistribute BGP 65001 metric 1 1 1 1 1
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Show ip route ospf | Inc O

O E2 5.0.0.0/8 [110/2297856] via 10.1.13.3, 00:01:19, Serial0/1
O E2    10.1.45.0 [110/1] via 10.1.13.3, 00:01:19, Serial0/1
```

On R5

```
R5#Show ip route eigrp | Inc D

D EX 1.0.0.0/8 [170/2560512256] via 10.1.45.4, 00:01:01, Serial0/1
D EX    10.1.13.0 [170/2560512256] via 10.1.45.4, 00:01:01, Serial0/1
```

On R6

```
R6#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 6.0.0.0          0.0.0.0                  0         32768 i
```

To configure simple VPN:

On R4

```
R4(config)#ip vrf CB
R4(config-vrf)#route-target Both 1:105
```

To verify the configuration:

On R6

```
R6#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path

*> 1.0.0.0          10.1.46.4                              0 65001 ?
*> 5.0.0.0          10.1.46.4                              0 65001 ?
*> 6.0.0.0          0.0.0.0                  0         32768 i
*> 10.1.13.0/24     10.1.46.4                              0 65001 ?
*> 10.1.45.0/24     10.1.46.4                              0 65001 ?

R6#Show ip route bgp | Inc B
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
B    1.0.0.0/8 [20/0] via 10.1.46.4, 00:01:22
B    5.0.0.0/8 [20/0] via 10.1.46.4, 00:01:22
B       10.1.13.0 [20/0] via 10.1.46.4, 00:01:22
B       10.1.45.0 [20/0] via 10.1.46.4, 00:01:22
```

On R1

```
R1#Show ip route ospf | Inc O

O E2 5.0.0.0/8 [110/2297856] via 10.1.13.3, 00:05:40, Serial0/1
O E2 6.0.0.0/8 [110/1] via 10.1.13.3, 00:02:03, Serial0/1
O E2    10.1.45.0 [110/1] via 10.1.13.3, 00:05:40, Serial0/1
```

On R5

```
R5#Show ip route eigrp | Inc D

D EX 1.0.0.0/8 [170/2560512256] via 10.1.45.4, 00:05:04, Serial0/1
D EX 6.0.0.0/8 [170/2560512256] via 10.1.45.4, 00:02:35, Serial0/1
D EX    10.1.13.0 [170/2560512256] via 10.1.45.4, 00:05:04, Serial0/1
```

Task 6

Reconfigure the Routers to support an environment where R5 can exchange routes with R1 and R6 but R6 cannot exchange routes with R1. This task may override the configuration of the previous task. Use any value to accomplish this task.

Step One:

Remove the route-target 1:105:

On R4

```
R4(config)#ip vrf CB
R4(config-vrf)#NO route-target b 1:105
```

Step Two:

Add a new route-target to VRF CB, this route-target should be configured to both import and export 1:506 between the two VRFs (CA and CB) on R4 so R5 and R6 can exchange routes:

```
R4(config-vrf)#route-target both 1:506

R4(config)#IP vrf CA
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-vrf)#route-target both 1:506
```

Note the following verifies that R1 and R5 share routes, because they both import and export the same route-target of 1:105; R5 and R6 share routes because they both import and export route-target 1:506; but since R1 and R6 do NOT import or export any route-targets amongst themselves, they will NOT share any routes.

On R1

```
R1#Show ip route ospf | Inc O

O E2 5.0.0.0/8 [110/2297856] via 10.1.13.3, 00:34:42, Serial0/1
O E2    10.1.45.0 [110/1] via 10.1.13.3, 00:34:42, Serial0/1
```

On R5

```
R5#Show ip route eigrp

D EX 1.0.0.0/8 [170/2560512256] via 10.1.45.4, 00:34:03, Serial0/1
D EX 6.0.0.0/8 [170/2560512256] via 10.1.45.4, 00:02:34, Serial0/1
     10.0.0.0/24 is subnetted, 2 subnets
D EX    10.1.13.0 [170/2560512256] via 10.1.45.4, 00:34:03, Serial0/1
```

On R6

```
R6#Show ip route bgp

B    5.0.0.0/8 [20/0] via 10.1.46.4, 00:03:55
     10.0.0.0/24 is subnetted, 2 subnets
B       10.1.45.0 [20/0] via 10.1.46.4, 00:03:55
```

Note the output of the above show command reveals that R5 participates in both VPNs:

On R4

```
R4#Show ip bgp vpnv4 vrf CA 5.0.0.0 | Inc RT

      Extended Community: RT:1:105 RT:1:506
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Task 7

The PE routers should be configured such that they allow R1 to manage R5 and R6 with a new loopback address. Neither R5 nor R6 should be able to see the loopback from each other as they are not managing each other.
Add the following loopback interfaces to the routers; you may override the previous configuration to accomplish this task: No other routes should be advertised.

R1- Loopback 100 100.100.11/32
R5- Loopback 100 100.100.55/32
R6- Loopback 100 100.100.66/32

On R1

```
R1(config)#int lo1
R1(config-if)#ip addr 100.100.100.11 255.255.255.255

R1(config-if)#router ospf 1
R1(config-router)#network 100.100.100.11 0.0.0.0 area 0
```

On R5

```
R5(config)#int lo1
R5(config-if)#ip addr 100.100.100.55 255.255.255.255

R5(config-if)#router eigrp 100
R5(config-router)#network 100.100.100.55 0.0.0.0
```

On R6

```
R6(config)#int lo1
R6(config-if)#ip addr 100.100.100.66 255.255.255.255

R6(config-if)#router bgp 65006
R6(config-router)#Network 100.100.100.66 mask 255.255.255.255
```

Configuring the Management VPN:

On R3

The following Prefix-list references Lo1 on R1:

```
R3(config)#ip prefix-list TST permit 100.100.100.11/32
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The following route-map references loopback 100 of R1 and assigns an extended community of RT 1:156, this is how a Route-target is assigned to a specific route:

```
R3(config)#Route-map Manage permit 10
R3(config-route-map)#match ip addr prefix TST
R3(config-route-map)#set extcommunity rt 1:156
R3(config)#Route-map Manage permit 90
```

The route-map from the previous step (Manage) is exported out, and route-target 1:56 is imported:

```
R3(config-route-map)#ip vrf CA
R3(config-vrf)#export map Manage
R3(config-vrf)#route-target import 1:56
```

On R4

```
R4(config)#ip prefix-list TST seq 5 permit 100.100.100.55/32
R4(config)#ip prefix-list TST seq 10 permit 100.100.100.66/32

R4(config)#route-map NMS permit 10
R4(config-route-map)#match ip addr prefix TST
R4(config-route-map)#set extcommunity rt 1:56
R4(config)#route-map NMS permit 90

R4(config)#ip vrf CA
R4(config-vrf)#export map NMS
R4(config-vrf)#route-target import 1:156
R4(config-vrf)#NO route-target b 1:506

R4(config)#ip vrf CB
R4(config-vrf)#export map NMS
R4(config-vrf)#route-target import 1:156
R4(config-vrf)#NO route-target b 1:506
```

To verify the configuration:

On R6

```
R6#Show ip route bgp | Inc B

B       100.100.100.11 [20/0] via 10.1.46.4, 00:03:28
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

On R5

```
R5#Sh ip rou eigrp | Inc D

D EX    100.100.100.11 [170/2560512256] via 10.1.45.4, 00:02:53, Serial0/1
D EX    10.1.13.0 [170/2560512256] via 10.1.45.4, 02:58:25, Serial0/1
```

On R1

```
R1#Show ip route ospf | Inc O

O E2    100.100.100.66 [110/1] via 10.1.13.3, 00:02:55, Serial0/1
O E2    100.100.100.55 [110/2297856] via 10.1.13.3, 02:18:28, Serial0/1
O E2 5.0.0.0/8 [110/2297856] via 10.1.13.3, 03:00:20, Serial0/1
O E2    10.1.45.0 [110/1] via 10.1.13.3, 03:00:20, Serial0/1
```

To test the configuration:

On R5

```
R5#Ping 100.100.100.11 sou lo1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.100.11, timeout is 2 seconds:
Packet sent with a source address of 100.100.100.55
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/56 ms
```

On R6

```
R6#Ping 100.100.100.11 source Lo1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.100.11, timeout is 2 seconds:
Packet sent with a source address of 100.100.100.66
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

On R1

```
R1#Ping 100.100.100.55 source Lo1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.100.55, timeout is 2 seconds:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Packet sent with a source address of 100.100.100.11
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/56 ms

R1#Ping 100.100.100.66 source Lo1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.100.66, timeout is 2 seconds:
Packet sent with a source address of 100.100.100.11
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms
```

Task 8

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

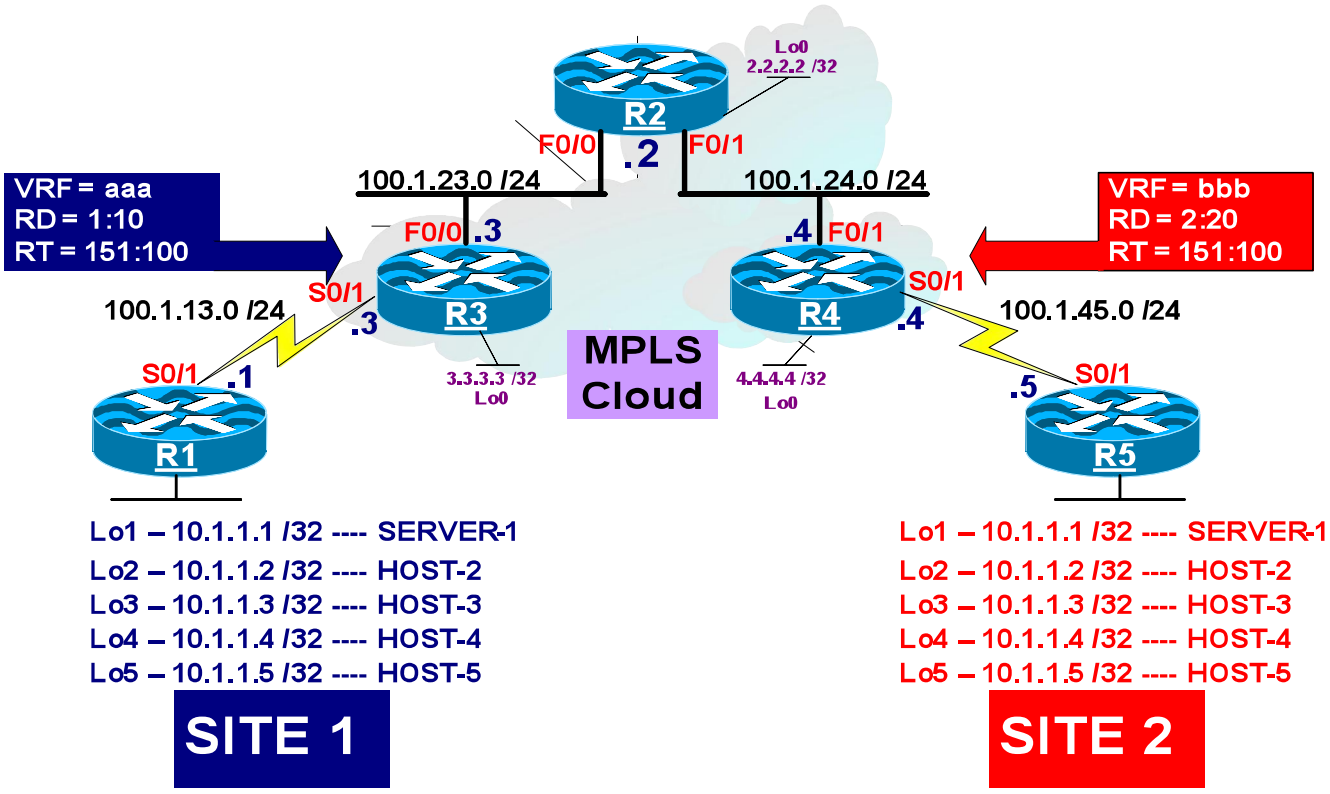
Paul Negron
CCIE #14856
SP

MPLS & NAT

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1 – MPLS and NAT - I



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure the Loopback interfaces of R1 and R5 based on the above diagram.
- Configure the F0/0 interface of R2 and R3 in VLAN 100.
- Configure the F0/1 interface of R2 and R4 in VLAN 200.
- Configure the rest of the routers according to the above diagram.

Task 1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Configure OSPF on the core routers (R2, R3 and R4); you should run OSPF area 0 on the F0/0 interfaces of R2 and R3, the F0/1 interfaces of R2 and R4, and the Loopback 0 interfaces of R2, R3 and R4. The CE rouers, R1 and R5 should be configured with a static default route pointing to their next hop router.

To configure the CE routers:

On R1

```
R1(config)#IP route 0.0.0.0 0.0.0.0 100.1.13.3
```

On R5

```
R5(config)#IP route 0.0.0.0 0.0.0.0 100.1.45.4
```

To configure the core routers:

On R2

```
R2(config)#Router ospf 1
R2(config-router)#Netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#Netw 100.1.23.2 0.0.0.0 area 0
R2(config-router)#Netw 100.1.24.2 0.0.0.0 area 0
```

On R3

```
R3(config)#Router ospf 1
R3(config-router)#Netw 100.1.23.3 0.0.0.0 area 0
R3(config-router)#Netw 3.3.3.3 0.0.0.0 area 0
```

On R4

```
R4(config)#Router ospf 1
R4(config-router)#Netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#Netw 100.1.24.4 0.0.0.0 area 0
```

To verify the configuration:

On R2

```
R2#Show ip ospf neighbor
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Neighbor ID     Pri   State           Dead Time   Address         Interface
4.4.4.4           1   FULL/BDR        00:00:33    100.1.24.4      FastEthernet0/1
3.3.3.3           1   FULL/BDR        00:00:33    100.1.23.3      FastEthernet0/0

R2#Show ip route ospf | Inc O

O        3.3.3.3 [110/2] via 100.1.23.3, 00:10:53, FastEthernet0/0
O        4.4.4.4 [110/2] via 100.1.24.4, 00:10:35, FastEthernet0/1
```

Task 2

Configure LDP between the core routers. These routers should use their Loopback0 interface as their LDP router-id.

On R2, R3 and R4

```
Rx(config)#Mpls label protocol ldp
Rx(config)#Mpls ldp router-id Lo0
```

On R3

```
R3(config)#Int F0/0
R3(config-if)#MPLS IP
```

On R2

```
R2(config)#Int F0/0
R2(config-if)#MPLS IP

R2(config-if)#Int F0/1
R2(config-if)#MPLS IP
```

On R4

```
R4(config)#Int F0/1
R4(config-if)#MPLS IP
```

To Verify the configuration:

On R2

```
R2#Show mpls ldp neighbor
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    Peer LDP Ident: 4.4.4.4:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 4.4.4.4.60890 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 9/10; Downstream
        Up time: 00:01:05
        LDP discovery sources:
          FastEthernet0/1, Src IP addr: 100.1.24.4
        Addresses bound to peer LDP Ident:
          100.1.24.4      100.1.45.4      4.4.4.4
    Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 2.2.2.2:0
        TCP connection: 3.3.3.3.18225 - 2.2.2.2.646
        State: Oper; Msgs sent/rcvd: 9/10; Downstream
        Up time: 00:01:00
        LDP discovery sources:
          FastEthernet0/0, Src IP addr: 100.1.23.3
        Addresses bound to peer LDP Ident:
          100.1.23.3      100.1.13.3      3.3.3.3
```

On R3

```
R3#Show mpls forwarding-table

Local  Outgoing    Prefix           Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id     switched   interface
16     Pop tag     2.2.2.2/32       0          Fa0/0       100.1.23.2
17     Pop tag     100.1.24.0/24    0          Fa0/0       100.1.23.2
18     17          4.4.4.4/32       0          Fa0/0       100.1.23.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 100. The ONLY BGP peering relationship should be VPNV4. These two neighbors should use their Lo0 interfaces for their peering.

On R3

```
R3(config)#Router bgp 100
R3(config-router)#Neighbor 4.4.4.4 remote-as 100
R3(config-router)#Neighbor 4.4.4.4 update-source Lo0

R3(config-router)#Address-family VPNV4 Unicast
R3(config-router-af)#Neighbor 4.4.4.4 Act
R3(config-router-af)#Neighbor 4.4.4.4 Send-community Ext
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4(config)#Router bgp 100
R4(config-router)#Neighbor 3.3.3.3 remote-as 100
R4(config-router)#Neighbor 3.3.3.3 update-source Lo0

R4(config-router)#Address-family VPNV4 Unicast
R4(config-router-af)#Neighbor 3.3.3.3 Act
R4(config-router-af)#Neighbor 3.3.3.3 Send-community Ext
```

To verify the configuration:

On R3

```
R3#Show ip bgp vpnv4 all Summary | B Neigh

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4   100       8       8        1    0    0 00:02:02        0
```

Task 4

Configure the following VRFs, RDs and route-targets on the PE routers:

Router	VRF Name	RD	Route-Target	Interface
R3	aaa	1:10	Route-target Both 151:100	S0/1
R4	bbb	2:20	Route-target Both 151:100	S0/1

On R3

```
R3(config)#IP VRF aaa
R3(config-vrf)#RD 1:10
R3(config-vrf)#Route-target Both 151:100

R3(config)#Int S0/1
R3(config-if)#IP VRF Forwarding aaa
R3(config-if)#IP address 100.1.13.3 255.255.255.0
```

On R4

```
R4(config)#IP VRF bbb
R4(config-vrf)#RD 2:20
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-vrf)#Route-target Both 151:100

R4(config)#Int S0/1
R4(config-if)#IP VRF Forwarding bbb
R4(config-if)#IP address 100.1.45.4 255.255.255.0
```

To verify the configuration:

On R3

```
R3#Show ip vrf detail

VRF aaa; default RD 1:10; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:151:100
  Import VPN route-target communities
    RT:151:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

On R4

```
R4#Show ip vrf detail

VRF bbb; default RD 2:20; default VPNID <not set>
  Interfaces:
    Se0/1
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:151:100
  Import VPN route-target communities
    RT:151:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Ensure that the hosts in Site-1 can access the server in Site 2 and vice versa. You should configure the CE routers (R1 and R5). Use the following translation chart:

Rouer	Inside Local	Inside Global
R1	**10.1.1.1** **10.1.1.2 – 10.1.1.5**	**1.1.1.1** **1.1.1.2 – 1.1.1.5**
R5	10.1.1.1 10.1.1.2 – 10.1.1.5	5.5.5.1 5.5.5.2 – 5.5.5.5

A static route for network 1.1.1.0 /24 is configured and redistributed into the vrf aaa on R3. This is done to provide reachability to the hosts connected to R5.

On R3

```
R3(config)#IP Route vrf aaa 1.1.1.0 255.255.255.0 100.1.13.1

R3(config)#Router bgp 100
R3(config-router)#Address-family IPv4 vrf aaa
R3(config-router-af)#Redistribute Static
R3(config-router-af)#Redistribute connected
```

The same is configured on R4:

On R4

```
R4(config)#IP Route vrf bbb 5.5.5.0 255.255.255.0 100.1.45.5

R4(config)#Router bgp 100
R4(config-router)#Address-family IPv4 vrf bbb
R4(config-router-af)#Redistribute Static
R4(config-router-af)#Redistribute Connected
```

To verify the configuration:

On R4

```
R4#Show ip route vrf bbb

Routing Table: bbb
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/24 is subnetted, 1 subnets
B       1.1.1.0 [200/0] via 3.3.3.3, 00:02:17
     100.0.0.0/24 is subnetted, 2 subnets
C       100.1.45.0 is directly connected, Serial0/1
B       100.1.13.0 [200/0] via 3.3.3.3, 00:02:17
     5.0.0.0/24 is subnetted, 1 subnets
S       5.5.5.0 [1/0] via 100.1.45.5
```

On R3

```
R3#Show ip route vrf aaa

Routing Table: aaa
Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

     1.0.0.0/24 is subnetted, 1 subnets
S       1.1.1.0 [1/0] via 100.1.13.1
     100.0.0.0/24 is subnetted, 2 subnets
B       100.1.45.0 [200/0] via 4.4.4.4, 00:02:06
C       100.1.13.0 is directly connected, Serial0/1
     5.0.0.0/24 is subnetted, 1 subnets
B       5.5.5.0 [200/0] via 4.4.4.4, 00:02:06
```

On R1

The NAT Inside and Outside interfaces are defined:

```
R1(config)#Int range Lo0 - 4
R1(config-if)#IP NAT Inside

R1(config)#Int S0/1
R1(config-if)#IP NAT Outside
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The following command translates the inside source IP address of 10.1.1.1 to 1.1.1.1 IP address:

```
R1(config)#IP NAT inside source static 10.1.1.1 1.1.1.1
```

An access-list is configured to identify the communication between inside sources with destination IP addresses:

```
R1(config)#Access-list 100 permit ip 10.1.1.0 0.0.0.255 5.5.5.0 0.0.0.255
```

The following configures a NAT pool that the inside hosts can use:

```
R1(config)#IP Nat pool TST 5.5.5.2 5.5.5.5 Prefix-length 24 type match-host
```

The last step is to configure the inside sources identified in ACL 100 to use the NAT pool called "TST":

```
R1(config)#IP NAT inside source list 100 pool TST
```

On R5

```
R5(config-if)#Int range Lo0 - 4
R5(config-if)#IP NAT Inside

R5(config)#Int S0/1
R5(config-if)#IP NAT Outside

R5(config)#IP NAT inside source static 10.1.1.1 5.5.5.1

R5(config)#Access-list 100 permit ip 10.1.1.0 0.0.0.255 1.1.1.0 0.0.0.255

R5(config)#IP Nat pool TST 1.1.1.2 1.1.1.5 Prefix-length 24
R5(config)#IP NAT inside source list 100 pool TST
```

To verify the configuration:

On R1

```
R1#Ping 5.5.5.1 source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/56/60 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Show ip nat translations

Pro Inside global      Inside local       Outside local      Outside global
icmp 1.1.1.1:2         10.1.1.1:2         5.5.5.1:2          5.5.5.1:2
--- 1.1.1.1            10.1.1.1           ---                ---

R1#Ping 5.5.5.1 Source Lo4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R1#Sh ip nat translation

Pro Inside global      Inside local       Outside local      Outside global
icmp 1.1.1.1:7         10.1.1.1:7         5.5.5.1:7          5.5.5.1:7
--- 1.1.1.1            10.1.1.1           ---                ---
icmp 1.1.1.5:8         10.1.1.5:8         5.5.5.5:8          5.5.5.5:8
--- 1.1.1.5            10.1.1.5           ---                ---
```

Task 6

Remove the configuration from the previous step and configure the PE routers to accomplish the same task.

On R1

```
R1(config)#Int range Lo0 - 4
R1(config-if-range)#NO IP NAT Inside

R1(config)#Int S0/1
R1(config-if)#NO IP NAT Outside

R1(config)#No Access-list 100
R1(config)#NO ip nat inside source static 10.1.1.1 1.1.1.1
R1(config)#NO ip nat inside source list 100 pool TST
R1(config)#NO ip nat pool TST 1.1.1.2 1.1.1.5 prefix-length 24
```

On R5

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config)#Int range Lo0 - 4
R5(config-if-range)#NO IP NAT Inside

R5(config)#Int S0/1
R5(config-if)#NO IP NAT Outside

R5(config)#NO access-list 100
R5(config)#NO ip nat inside source static 10.1.1.1 5.5.5.1
R5(config)#NO ip nat inside source list 100 pool TST
R5(config)#NO ip nat pool TST 1.1.1.2 1.1.1.5 prefix-length 24
```

NOTE: The configuration on the PE is identical to the configuration that was performed on the CEs with one difference; on the PEs the VRF MUST be referenced.

On R3

The inside and outside interfaces are defined; the interface facing the CE MUST be defined as inside, and the interface facing the core must be defined as outside.

```
R3(config)#Int S0/1
R3(config-if)#IP NAT Inside

R3(config)#Int F0/0
R3(config-if)#IP NAT Outside
```

A Static NAT is configured to translate any traffic with a source IP address of 10.1.1.1 to 1.1.1.1 IP address IN VRF "aaa":

```
R3(config)#IP NAT inside source static 10.1.1.1 1.1.1.1 vrf aaa
```

An access-list is configured to identify the communication between inside sources with destination IP addresses:

```
R3(config)#Access-list 100 permit ip 10.1.1.0 0.0.0.255 5.5.5.0 0.0.0.255
```

A NAT pool called "TST" is configured:

```
R3(config)#IP NAT Pool TST 1.1.1.2 1.1.1.5 Prefix-length 24 Type match-host
```

The last step is to configure the inside sources identified in ACL 100 to use the NAT pool called "TST" for VRF "aaa":

```
R3(config)#IP NAT inside source list 100 pool TST vrf aaa
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#Int S0/1
R4(config-if)#IP NAT Inside

R4(config)#Int F0/1
R4(config-if)#IP NAT Outside

R4(config)#IP NAT inside source static 10.1.1.1 5.5.5.1 vrf bbb

R4(config)#Access-list 100 permit ip 10.1.1.0 0.0.0.255 1.1.1.0 0.0.0.255

R4(config)#IP NAT Pool TST 5.5.5.2 5.5.5.5 prefix-length 24 type match-host

R4(config)#IP NAT Inside source list 100 pool TST vrf bbb
```

To verify the configuration:

On R3

```
R3#Show ip nat translations vrf aaa

Pro Inside global      Inside local       Outside local      Outside global
--- 1.1.1.1            10.1.1.1           ---                ---
```

To test the configuration:

On R1

```
R1#Ping 5.5.5.1 Source Lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```

On R3

```
R3#Show ip nat translation vrf aaa

Pro Inside global      Inside local       Outside local      Outside global
icmp 1.1.1.1:0         10.1.1.1:0         5.5.5.1:0          5.5.5.1:0
--- 1.1.1.1            10.1.1.1           ---                ---

R1#Ping 5.5.5.1 Source Lo4

Type escape sequence to abort.
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Sending 5, 100-byte ICMP Echos to 5.5.5.1, timeout is 2 seconds:
Packet sent with a source address of 10.1.1.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/56 ms
```

On R3

```
R3#Show ip nat translation vrf aaa

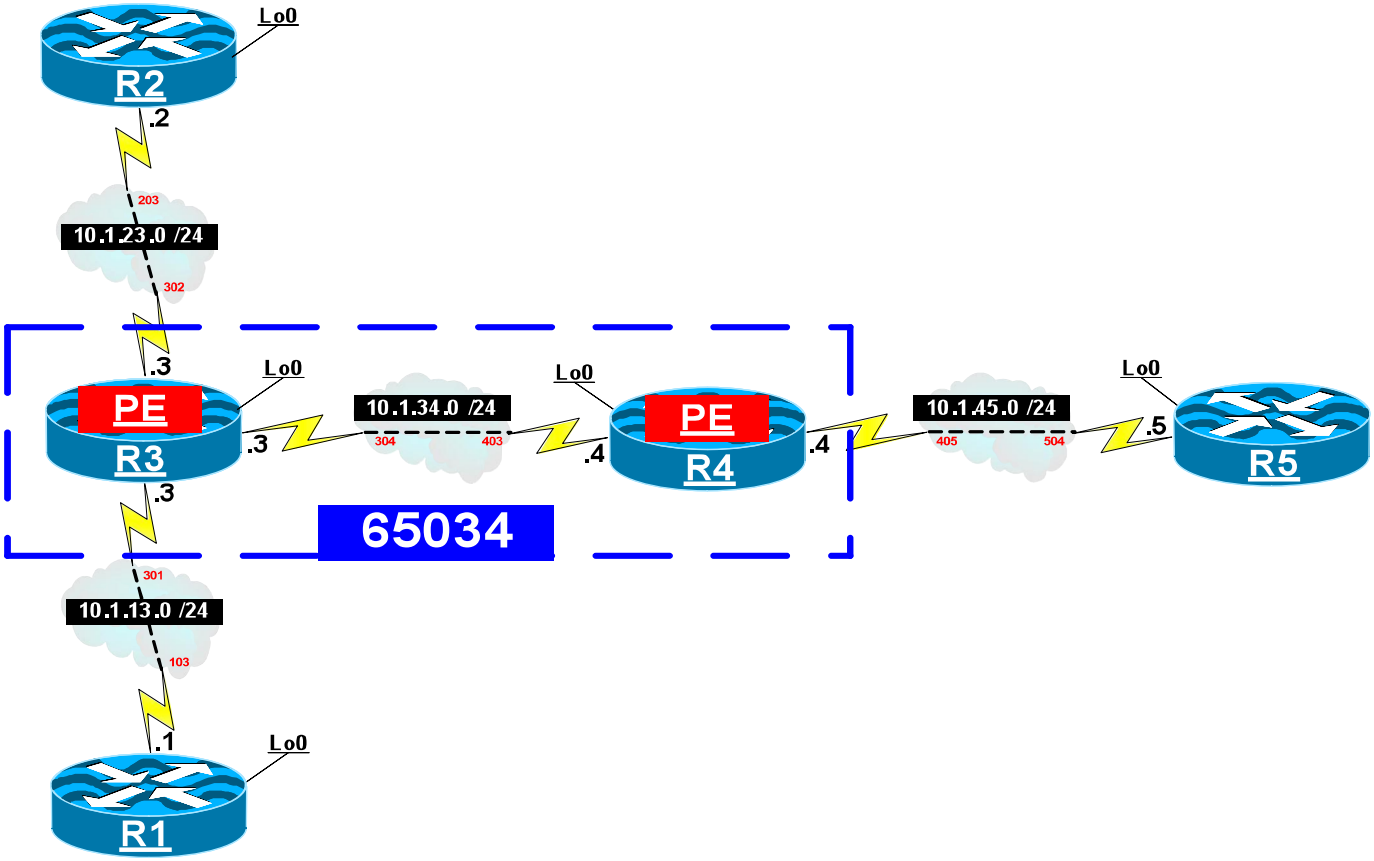
Pro Inside global      Inside local       Outside local      Outside global
icmp 1.1.1.1:2         10.1.1.1:2         10.1.1.2:2         10.1.1.2:2
--- 1.1.1.1            10.1.1.1           ---                ---
icmp 1.1.1.5:1         10.1.1.5:1         5.5.5.5:1          5.5.5.5:1
--- 1.1.1.5            10.1.1.5           ---                ---
```

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

LAB 4- MPLS and NAT - II



IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12	1.1.1.1 /32 10.1.12.1 /24
R2	Lo0 S0/0.21	1.1.1.1/ 32 10.1.23.2 /24
R3	Lo0 S0/0.31 S0/0.32	3.3.3.3/ 32 10.1.13.3/ 24 10.1.23.3/ 24
R4	Lo0 S0/0.43 S0/0.45	4.4.4.4 /32 10.1.34.4/ 24 10.1.45.4/ 24
R5	Lo0 S0/0.54	5.5.5.5 /32 10.1.45.5 /24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab Setup:

- Configure OSPF between Core routers R3 and R4 using process "1"

Task 1

Configure Core Routers (R3 and R4) to support MPLS VPN's using AS 65034.

On R3

```
R3(config)#mpls label prot ldp
R3(config)#mpls ldp router-id lo0

R3(config)#int s0/0.34
R3(config-subif)#mpls ip

R3(config-subif)#router bgp 65034
R3(config-router)#NO bgp def ipv4
R3(config-router)#neighbor 4.4.4.4 remot 65034
R3(config-router)#neighbor 4.4.4.4 update lo0

R3(config-router)#address vpnv4
R3(config-router-af)#neighbor 4.4.4.4 act
R3(config-router-af)#neighbor 4.4.4.4 send ext
R3(config-router-af)#neighbor 4.4.4.4 next-hop-s
```

On R4

```
R3(config)#mpls label prot ldp
R3(config)#mpls ldp router-id lo0

R3(config)#int s0/0.43
R3(config-subif)#mpls ip

R4(config-subif)#router bgp 65034
R4(config-router)#NO bgp def ipv4
R4(config-router)#neighbor 3.3.3.3 remot 65034
R4(config-router)#neighbor 3.3.3.3 update lo0

R4(config-router)#address vpnv4

R4(config-router-af)#neighbor 3.3.3.3 act
R4(config-router-af)#neighbor 3.3.3.3 send ext
R4(config-router-af)#neighbor 3.3.3.3 next-hop-s
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Create three VRF's for R1, R2 and R5 using info in the Table below:

VRF	RD	Route-Target	Interface
NAT-R1	3.3.3.3:100	65034:10	S0/0.31
NAT-R2	3.3.3.3:200	65034:20	S0/0.32
SERVICE	4.4.4.4:100	65034:30	S0/0.45

On R3

```
R3(config)#ip vrf NAT-R1
R3(config-vrf)#rd 3.3.3.3:100
R3(config-vrf)#route b 65034:10

R3(config)#ip vrf NAT-R2
R3(config-vrf)#rd 3.3.3.3:200
R3(config-vrf)#route b 65034:20

R3(config)#int s0/0.31
R3(config-subif)#ip vrf forwarding NAT-R1
R3(config-subif)#ip address 10.1.13.3 255.255.255.0

R3(config-subif)#int s0/0.32
R3(config-subif)#ip vrf forwarding NAT-R2
R3(config-subif)#ip address 10.1.23.3 255.255.255.0

R3#Ping vrf NAT-R1 10.1.13.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/8/20 ms

R3#Ping vrf NAT-R2 10.1.23.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.23.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/24/44 ms
```

On R4

```
R4(config)#ip vrf SERVICE
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-vrf)#rd 4.4.4.4:100
R4(config-vrf)#route b 65034:30

R4(config-vrf)#int s0/0.45
R4(config-subif)#ip vrf forwarding SERVICE
R4(config-subif)#ip address 10.1.45.4 255.255.255.0

R4#Ping vrf SERVICE 10.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.45.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/10/16 ms
```

Task 3

Configure EIGRP as the routing protocol between the PE and CE routers.

On R1

```
R1(config)#router eigrp 1
R1(config-router)#netw 10.1.13.0 0.0.0.255
R1(config-router)#netw 1.0.0.0
R1(config-router)#NO auto
```

On R2

```
R2(config)#router eigrp 1
R2(config-router)#netw 10.1.23.0 0.0.0.255
R2(config-router)#netw 1.0.0.0
R2(config-router)#NO auto
```

On R5

```
R5(config)#router eigrp 1
R5(config-router)#netw 10.1.45.0 0.0.0.255
R5(config-router)#NO auto
```

On R3

```
R3(config)#router eigrp 1
R3(config-router)#address ipv4 vrf NAT-R1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router-af)#NO auto
R3(config-router-af)#autonomous-system 1
R3(config-router-af)#netw 10.1.13.0 0.0.0.255
R3(config-router-af)#redistribute bgp 65034 metric 1 1 1 1 1

R3(config-router-af)#address ipv4 vrf NAT-R2
R3(config-router-af)#NO auto
R3(config-router-af)#autonomous-system 1
R3(config-router-af)#netw 10.1.23.0 0.0.0.255
R3(config-router-af)#redistribute bgp 65034 metric 1 1 1 1 1

R3(config-router-af)#router bgp 65034
R3(config-router)#address ipv4 vrf NAT-R1
R3(config-router-af)#redistribute eigrp 1

R3(config-router-af)#address ipv4 vrf NAT-R2
R3(config-router-af)#redistribute eigrp 1

R3#Sh ip route vrf NAT-R1 | b Gateway

Gateway of last resort is not set

     1.0.0.0/32 is subnetted, 1 subnets
D       1.1.1.1 [90/2297856] via 10.1.13.1, 00:03:50, Serial0/0.31
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/0.31

R3#Sh ip route vrf NAT-R2 | b Gateway

Gateway of last resort is not set

     1.0.0.0/32 is subnetted, 1 subnets
D       1.1.1.1 [90/2297856] via 10.1.23.2, 00:02:12, Serial0/0.32
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.23.0 is directly connected, Serial0/0.32

R3#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 3.3.3.3:100 (default for vrf NAT-R1)
*> 1.1.1.1/32       10.1.13.1          2297856         32768 ?
*> 10.1.13.0/24     0.0.0.0                  0         32768 ?
Route Distinguisher: 3.3.3.3:200 (default for vrf NAT-R2)
*> 1.1.1.1/32       10.1.23.2          2297856         32768 ?
*> 10.1.23.0/24     0.0.0.0                  0         32768 ?
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Notice how both 1.1.1.1 routes are maintained in the BGP table. They both have unique RD information. This will allow the prefixes to be routed through the Core with no problems.

On R4

```
R4(config)#router eigrp 1
R4(config-router)#address ipv4 vrf SERVICE
R4(config-router-af)#NO auto
R4(config-router-af)#autonomous-sys 1
R4(config-router-af)#netw 10.1.45.0 0.0.0.255


R4#Sh ip route vrf SERVICE | b Gatew

Gateway of last resort is not set

     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/0.45
```

The 5.5.5.5 prefix is NOT included with the update. This prefix will be statically configured to help support the NAT configuration.

Task 4

Provide Internet Access for R1's loopback 0 interface without using a dynamic routing protocol between the customer and the R2 PE Router.

On R4

```
R4(config)#ip nat pool FILESERVER 192.168.1.0 192.168.1.254 prefix-length 24
```

Very standard NAT pool configuration.

```
R4(config)#ip route vrf SERVICE 192.168.1.0 255.255.255.0 null 0
```

This route is added to enable the CE router to route back to the translated addresses.

```
R4(config)#router eigrp 1
R4(config-router)#address ipv4 vrf SERVICE
R4(config-router-af)#redistribute static
```

On R5

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5#Sh ip route eigrp

D EX 192.168.1.0/24 [170/2169856] via 10.1.45.4, 00:00:18, Serial0/0.54
```

The CE router should be able to reach ANY translated address.

On R4

```
R4#Sh ip bgp vpnv4 all
```

R4 does not have ANY VPN routes. The reason for this is that the SERVICE VPN is not redistributing routes. Also, the routes from the PE router are not seen due to the default route-target filter that is enabled on a PE that does NOT have the route targets configured.

There are 2 ways to see the routes:
1) Remove the route-target filter with the "no bgp default route-target filter" command.
2) Configure the VRF's with the correct route targets and do not apply it anywhere.

```
R4(config)#ip vrf NAT-R1
R4(config-vrf)#rd 3.3.3.3:100
R4(config-vrf)#route 65034:10

R4(config)#ip vrf NAT-R2
R4(config-vrf)#rd 3.3.3.3:200
R4(config-vrf)#route 65034:20

R4#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 3.3.3.3:100 (default for vrf NAT-R1)
*>i1.1.1.1/32       3.3.3.3            2297856    100      0 ?
*>i10.1.13.0/24     3.3.3.3                  0    100      0 ?
Route Distinguisher: 3.3.3.3:200 (default for vrf NAT-R2)
*>i1.1.1.1/32       3.3.3.3            2297856    100      0 ?
*>i10.1.23.0/24     3.3.3.3                  0    100      0 ?

R4(config)#ip route vrf NAT-R1 5.0.0.0 255.0.0.0 s0/0.45 10.1.45.5

R4(config)#ip route vrf NAT-R2 5.0.0.0 255.0.0.0 s0/0.45 10.1.45.5
```

The second option was chosen to see the routes. We need to configure a static route for each VRF to be able to reach the common service for every VPN. We cannot configure the static routes until the VRF exists.

```
R4(config)#router bgp 65034
R4(config-router)#address ipv4 vrf NAT-R1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router-af)#redistribute static
R4(config-router-af)#address ipv4 vrf NAT-R2
R4(config-router-af)#redistribute static
```

On R1

```
R1#Sh ip route eigrp

D EX 5.0.0.0/8 [170/2560512256] via 10.1.13.3, 00:00:17, Serial0/0.12
```

On R2

```
R2#Sh ip route eigrp

D EX 5.0.0.0/8 [170/2560512256] via 10.1.23.3, 00:00:32, Serial0/0.23
```

On R4

```
R4(config)#int s0/0.45
R4(config-subif)#ip nat outside

R4(config-subif)#int s0/0.43
R4(config-subif)#ip nat inside
```

Be sure to label the inside and inside interfaces for NAT. In this case, the interface pointing to the SERVICE router is the outside but ANY VRF interface can be an inside or outside interface depending on the application of NAT.

```
R4(config)#access-list 111 permit ip any 5.0.0.0 0.255.255.255
R4(config)#route-map NAT-PE permit 10
R4(config-route-map)#match ip address 111
```

The acess-list is applied to a route-map to allow multiple translation destinations for complex NAT designs. We could have applied the access-list directly to NAT.

```
R4(config)#ip nat inside source route-map NAT-PE pool FILESERVER vrf NAT-R1

R4(config)#ip nat inside source route-map NAT-PE pool FILESERVER vrf NAT-R2
```

As you can see, we can provide different pools for different VRF's. However, by using a single pool for multiple VRF;s, we can ensure that the overlap will always seem unique to the server that is receiving these requests.

```
R1#Ping 5.5.5.5 so lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/56/116 ms

R2#Ping 5.5.5.5 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 24/47/92 ms

R4#Sh ip nat translations

Pro Inside global      Inside local       Outside local      Outside global
icmp 192.168.1.1:1     1.1.1.1:1          5.5.5.5:1          5.5.5.5:1
icmp 192.168.1.2:0     1.1.1.1:0          5.5.5.5:0          5.5.5.5:0
```

All of the pings work and both translatons can be seen on R4.

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

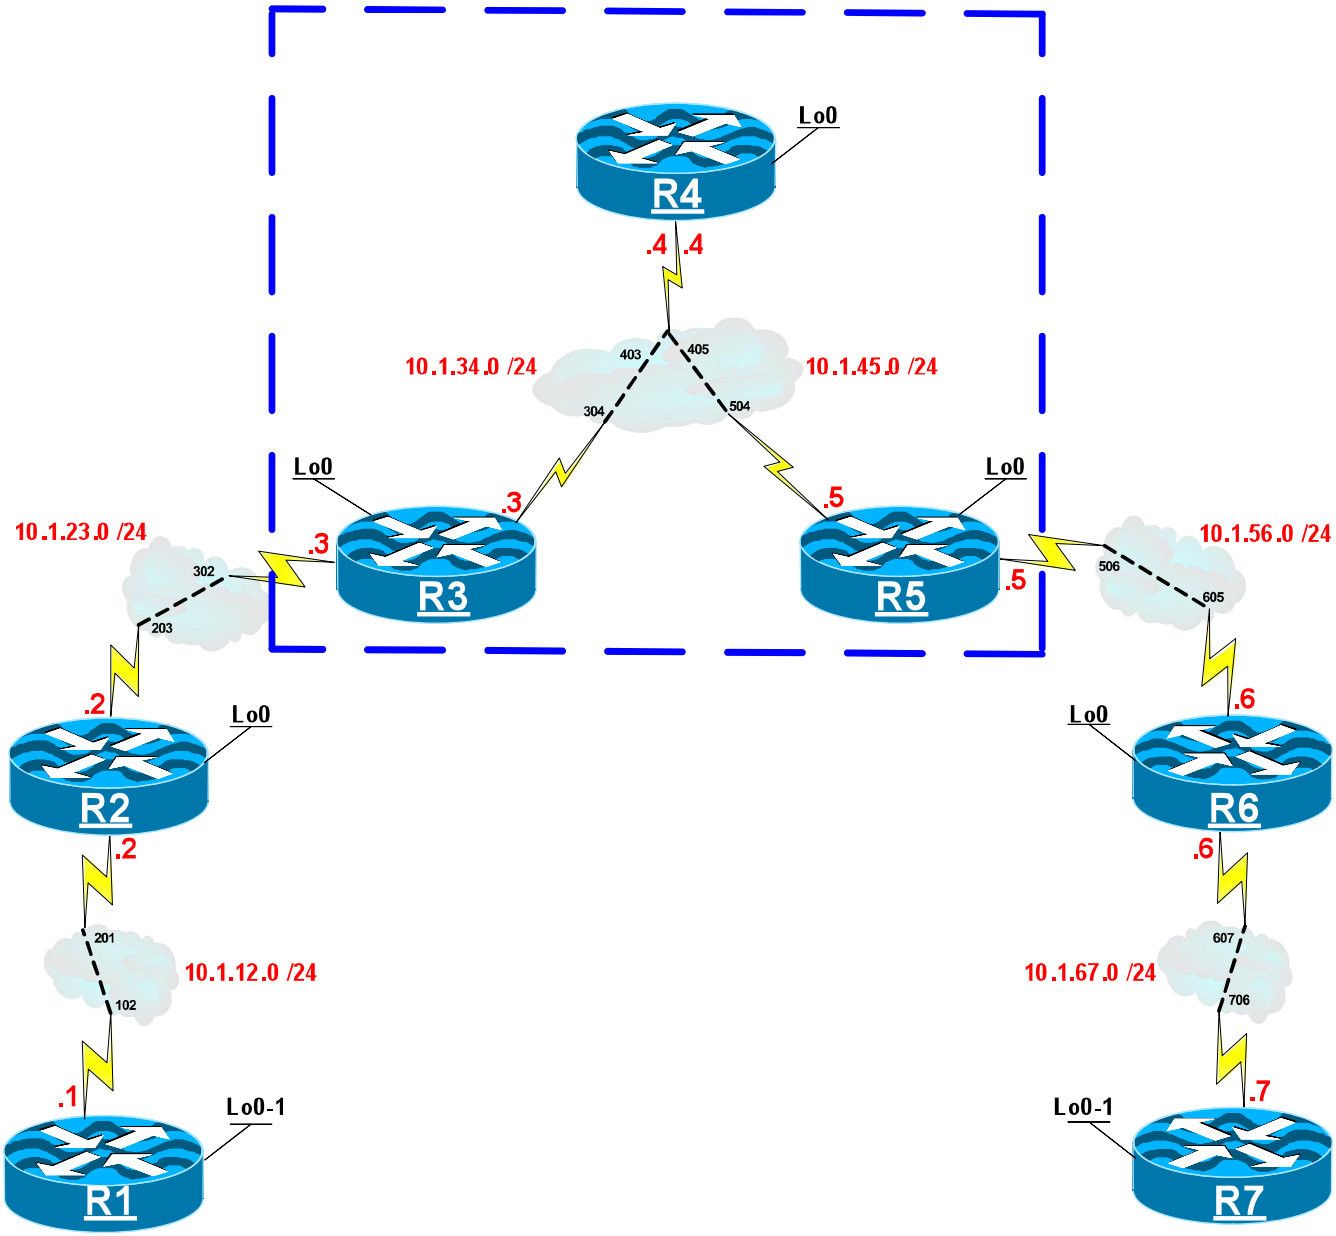
Paul Negron
CCIE #14856
SP

Carrier Supporting Carrier

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

LAB 1
Carrier Supporting Carrier (CSC)



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 Lo1 S0/0.12	1.1.1.1 /32 11.11.11.1/ 32 10.1.12.1 /24
R2	Lo0 S0/0.21 S1/0.23	2.2.2.2 /32 10.1.12.2 /24 10.1.23.2/ 24
R3	Lo0 S1/0.32 S1/0.34	3.3.3.3 /32 10.1.23.3/ 24 10.1.34.3/ 24
R4	Lo0 S0/0.43 S0/0.45	4.4.4.4 /32 10.1.34.4/ 24 10.1.45.4 /24
R5	Lo0 S0/0.54 S0/0.56	5.5.5.5 /32 10.1.45.5/ 24 10.1.56.5/ 24
R6	Lo0 S1/0.65 S1/0.67	6.6.6.6 /32 10.1.56.6/ /24 10.1.67.6 /24
R7	Lo0 Lo1 S1/0.76	7.7.7.7 /32 77.77.77.7 /32 10.1.67.7/ 32

Task 1

Configure EIGRP on Customer Routers R1,R2,R6 and R7 using an AS of "1". Core routers R3,R4 and R5 with an OSPF process of "1". Announce all loopbacks and interconnecting interfaces.

On R1

```
R1(config)#router eigrp 1
R1(config-router)#NO auto
R1(config-router)#net 1.0.0.0
R1(config-router)#net 11.0.0.0
R1(config-router)#net 10.0.0.0
R1(config-router)#pass lo0
```

On R2

```
R2(config)#router eigrp 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-router)#NO auto
R2(config-router)#netw 10.0.0.0
R2(config-router)#netw 2.0.0.0
R2(config-router)#pass lo0
```

On R6

```
R6(config)#router eigrp 1
R6(config-router)#NO auto
R6(config-router)#netw 10.0.0.0
R6(config-router)#net 6.0.0.0
R6(config-router)#passive-interface lo0
```

On R7

```
R7(config)#router eigrp 1
R7(config-router)#NO auto
R7(config-router)#netw 10.0.0.0
R7(config-router)#netw 7.0.0.0
R7(config-router)#netw 77.0.0.0
R7(config-router)#pass lo0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.0.0 0.0.255.255 area 0
R3(config-router)#net 3.3.3.3 0.0.0.0 area 0
R3(config-router)#passive-interface lo0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#network 10.1.0.0 0.0.255.255 area 0
R4(config-router)#network 4.4.4.4 0.0.0.0 area 0
R4(config-router)#passive lo0
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#network 10.1.0.0 0.0.255.255 area 0
R5(config-router)#network 5.5.5.5 0.0.0.0 area 0
R5(config-router)#passive lo0
```

Verify the Configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Sh ip ei n

IP-EIGRP neighbors for process 1
H   Address                 Interface       Hold Uptime   SRTT   RTO  Q  Seq
                                            (sec)         (ms)       Cnt Num
0   10.1.12.2               Se0/0.12          14 00:06:46   15    200  0  5
```

On R2

```
R2#Sh ip route | i D

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route
D       1.1.1.1 [90/2297856] via 10.1.12.1, 00:07:20, Serial0/0.21
D       11.11.11.1 [90/2297856] via 10.1.12.1, 00:06:46, Serial0/0.21
```

On R3

```
R3#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O       4.4.4.4 [110/65] via 10.1.34.4, 00:01:35, Serial1/0.34
O       5.5.5.5 [110/129] via 10.1.34.4, 00:01:35, Serial1/0.34
O       10.1.45.0 [110/128] via 10.1.34.4, 00:01:35, Serial1/0.34
O       10.1.56.0 [110/192] via 10.1.34.4, 00:01:35, Serial1/0.34
```

On R4

```
R4#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O       3.3.3.3 [110/65] via 10.1.34.3, 00:01:40, Serial0/0.43
O       5.5.5.5 [110/65] via 10.1.45.5, 00:01:40, Serial0/0.45
O       10.1.23.0 [110/128] via 10.1.34.3, 00:01:40, Serial0/0.43
O       10.1.56.0 [110/128] via 10.1.45.5, 00:01:40, Serial0/0.45
```

On R5

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O       3.3.3.3 [110/129] via 10.1.45.4, 00:01:35, Serial1/0.54
O       4.4.4.4 [110/65] via 10.1.45.4, 00:01:35, Serial1/0.54
O       10.1.23.0 [110/192] via 10.1.45.4, 00:01:35, Serial1/0.54
O       10.1.34.0 [110/128] via 10.1.45.4, 00:01:35, Serial1/0.54
```

On R6

```
R6#Sh ip route | i D

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route
D       7.7.7.7 [90/2297856] via 10.1.67.7, 00:05:58, Serial0/0.67
D       77.77.77.7 [90/2297856] via 10.1.67.7, 00:05:53, Serial0/0.67
```

On R7

```
R7#Sh ip ei n

IP-EIGRP neighbors for process 1
H   Address                 Interface       Hold  Uptime    SRTT   RTO  Q   Seq
                                            (sec)           (ms)       Cnt Num
0   10.1.67.6               Se0/0.76          10  00:06:13   444  2664  0   9
```

Task 2

Configure MP-BGP between the PE routers. The Core should only support VPNv4 routes. Enable LDP protocol between Provider devices.

On R3

```
R3(config)#mpls label range 300 399
R3(config)#mpls label protocol ldp

R3(config)#mpls ldp router-id loopback 0

R3(config)#int s1/0.34
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-subif)#mpls ip

R3(config)#router bgp 65000
R3(config-router)#NO bgp default ipv4-unicast
R3(config-router)#neighbor 5.5.5.5 remot 65000
R3(config-router)#neighbor 5.5.5.5 update-source loopback 0

R3(config-router)#address vpnv4 unicast
R3(config-router-af)#neighbor 5.5.5.5 activate
R3(config-router-af)#neighbor 5.5.5.5 send-community both
R3(config-router-af)#neighbor 5.5.5.5 next-hop-self
```

The MPLS range is not required but is used for clarity when confirming configurations during the lab exercises.

On R4

```
R4(config)#mpls label range 400 499
R4(config)#mpls label protocol ldp
R4(config)#mpls ldp router-id loopback 0

R4(config)#int s0/0.43
R4(config-subif)#mpls ip

R4(config-subif)#int s0/0.45
R4(config-subif)#mpls ip
```

On R5

```
R5(config)#mpls label range 500 599
R5(config)#mpls label protocol ldp
R5(config)#mpls ldp router-id loopback 0

R5(config)#interface s1/0.54
R5(config-if)#mpls ip

R5(config)#router bgp 65000
R5(config-router)#NO bgp default ipv4-unicast
R5(config-router)#neighbor 3.3.3.3 remot 65000
R5(config-router)#neighbor 3.3.3.3 update-source loopback 0

R5(config-router)#address vpnv4 unicast

R5(config-router-af)#neighbor 3.3.3.3 activate
R5(config-router-af)#neighbor 3.3.3.3 send-community both
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-router-af)#neighbor 3.3.3.3 next-hop-self
```

Verify the Configuration:

On R3

```
R3#Sh ip bgp vpnv4 all s

BGP router identifier 3.3.3.3, local AS number 65000
BGP table version is 1, main routing table version 1

Neighbor   V      AS   MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
5.5.5.5    4   65000         7          7        1    0    0 00:03:12        0
```

On R4

```
R4#Sh mpls int

Interface              IP             Tunnel      Operational
Serial0/0.43           Yes (ldp)      No          Yes
Serial0/0.45           Yes (ldp)      No          Yes
```

On R5

```
R3#Sh ip bgp vpnv4 all s

BGP router identifier 3.3.3.3, local AS number 65000
BGP table version is 1, main routing table version 1

Neighbor   V      AS   MsgRcvd   MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3    4   65000        12         12        1    0    0 00:08:24        0
```

Task 3

Add the VRF context to the PE routers for the Customers. The VRF name should be TST, the RD should be 200:200 for R3 and 300:300 for R5 with an RT of 1:10 on both routers. Complete the routing context to each customer site.

Adding the VRF Context:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#ip vrf TST
R3(config-vrf)#rd 200:200
R3(config-vrf)#route-target both 1:10

R3(config)#int s1/0.32
R3(config-subif)#ip vrf forwarding TST
R3(config-subif)#ip address 10.1.23.3 255.255.255.0

R3(config-subif)#router eigrp 1
R3(config-router)#address ipv4 vrf TST
R3(config-router-af)#autonomous-system 1
R3(config-router-af)#NO auto
R3(config-router-af)#network 10.1.23.3 0.0.0.0
R3(config-router-af)#redistribute bgp 65000 metric 1 1 1 1 1

R3(config-router-af)#router bgp 65000
R3(config-router)#address ipv4 vrf TST
R3(config-router-af)#redistribute eigrp 1
```

On R5

```
R5(config)#ip vrf TST
R5(config-vrf)#rd 300:300
R5(config-vrf)#route-target both 1:10

R5(config)#int s1/0.56
R5(config-subif)#ip vrf forwarding TST
R5(config-subif)#ip address 10.1.56.5 255.255.255.0

R5(config-subif)#router eigrp 1
R5(config-router)#address ipv4 vrf TST
R5(config-router-af)#autonomous-system 1
R5(config-router-af)#NO auto
R5(config-router-af)#network 10.1.56.5 0.0.0.0
R5(config-router-af)#redistribute bgp 65000 metric 1 1 1 1 1

R5(config-router-af)#router bgp 65000
R5(config-router)#address ipv4 vrf TST
R5(config-router-af)#redistribute eigrp 1
```

To verify the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Sh ip route eigrp | in D

D       2.2.2.2 [90/2297856] via 10.1.12.2, 00:43:58, Serial0/0.12
D       6.6.6.6 [90/3321856] via 10.1.12.2, 00:00:26, Serial0/0.12
D       7.7.7.7 [90/3833856] via 10.1.12.2, 00:00:26, Serial0/0.12
D       77.77.77.0 [90/3833856] via 10.1.12.2, 00:00:26, Serial0/0.12
D       10.1.23.0 [90/2681856] via 10.1.12.2, 00:43:59, Serial0/0.12
D       10.1.56.0 [90/3193856] via 10.1.12.2, 00:00:26, Serial0/0.12
D       10.1.67.0 [90/3705856] via 10.1.12.2, 00:00:26, Serial0/0.12

R1#Ping 7.7.7.7 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 7.7.7.7, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/67/88 ms
```

Task 4

Configure the CE routers to support MPLS LSP on the necessary interfaces. Allow the MPLS domain from CE to CE to flow through the Service Provider Network. Only an LDP Tunnel is permitted to be used.

Configure MPLS LDP on CE Routers:

On R1

```
R1(config)#mpls label range 100 199
R1(config)#mpls label protocol ldp
R1(config)#mpls ldp router-id loopback 0

R1(config)#int s0/0.12
R1(config-subif)#mpls ip
```

On R2

```
R2(config)#mpls label range 200 299
R2(config)#mpls label protocol ldp
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#mpls ldp router-id loopback 0

R2(config)#int s0/0.21
R2(config-subif)#mpls ip

R2(config-subif)#int s0/0.23
R2(config-subif)#mpls ip
```

The interface facing the PE is normally not labeled. The provider supports labeled interfaces toward the edge to allow the customer to participate in the MPLS service.

On R6

```
R6(config)#mpls label range 600 699
R6(config)#mpls label protocol ldp
R6(config)#mpls ldp router-id loopback 0

R6(config)#int s0/0.65
R6(config-subif)#mpls ip

R6(config-subif)#int s0/0.67
R6(config-subif)#mpls ip
```

On R7

```
R7(config)#mpls label range 700 799
R7(config)#mpls label protocol ldp
R7(config)#mpls ldp router-id loopback 0

R7(config)#int s0/0.76
R7(config-subif)#mpls ip
```

On R3

```
R3(config)#int s1/0.32
R3(config-subif)#mpls ip
```

***Feb 2 10:19:51.915: %SYS-5-CONFIG_I: Configured from console by console**

***Feb 2 10:19:52.227: %LDP-5-NBRCHG: LDP Neighbor (vrf TST) 2.2.2.2:0 (2) is UP**

This is the configuration that is not usually recommended but is required to allow the MPLS service to work. The output also reveals that this particular LDP session is within the VRF.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R5

```
R5(config)#int s1/0.56
R5(config-subif)#mpls ip
```

To verify the configuration:

On R7

```
R7#Trace 1.1.1.1

Type escape sequence to abort.
Tracing the route to 1.1.1.1

  1 10.1.67.6 [MPLS: Label 602 Exp 0] 100 msec 188 msec 108 msec
  2 10.1.45.5 [MPLS: Label 508 Exp 0] 148 msec 180 msec 116 msec
  3 10.1.45.4 [MPLS: Labels 401/303 Exp 0] 148 msec 48 msec 156 msec
  4 10.1.23.3 [MPLS: Label 303 Exp 0] 92 msec 80 msec 116 msec
  5 10.1.23.2 [MPLS: Label 200 Exp 0] 60 msec 100 msec 124 msec
  6 10.1.12.1 80 msec *  88 msec
```

The labels in your output may differ slightly than the output above. We first start with a trace from R7 to R1. Each line is explained below and is used as initial understanding for the additional output that follows:

Line 1 – Label 602 is imposed on R7.

Line 2 – R6 expects to see label "602" which is local and switches it with the label 508, learned from R5.

Line-3- R5 is the PE Router and will see the VPNv4 label from R3 as well as the label from R4 for Unicast Forwaring. Label "401" represents the "top most label" and will be imposed on the egress interface s1/0.54. Label "303" is the VPNv4 label that was imposed on interface s1/0.56 ingress.

Line 4- R4 receives label "401" and maps the label to a "pop tag" action. When it reaches R3 it has the "top most label" removed.

Line 5- R3 receives label "303" and maps the label to "200" which was received by R2.

Line 6- The label is popped off on R2. There is no further label processing.

```
R7#Sh mpls forwarding-table | i 1.1.1.1

702     602          1.1.1.1/32        0           Se0/0.76    point2point
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R7#Sh ip cef 1.1.1.1 detail | b via

  via 10.1.67.6, Serial0/0.76, 0 dependencies
    next hop 10.1.67.6, Serial0/0.76
    valid cached adjacency
    tag rewrite with Se0/0.76, point2point, tags imposed: {602}
```

The FIB and LFIB are shown above but the FIB will be used to impose the "602" label.

On R6

```
R6#Sh mpls ldp bindings 1.1.1.1 32

  tib entry: 1.1.1.1/32, rev 12
     local binding:  tag: 602
     remote binding: tsr: 10.1.56.5:0, tag: 508
     remote binding: tsr: 7.7.7.7:0, tag: 702

R6#Sh mpls forwarding-table | i 1.1.1.1

602     508         1.1.1.1/32        0          Se0/0.65   point2point
```

In this case, since the incoming packet has been labeled on the egress of R7, R6 receives the "602" label and swaps it with "508".

On R5

```
R5#Sh mpls forwarding-table | i 1.1.1.1

508     401         1.1.1.1/32[V]     1580          Se1/0.54   point2point
```

At first glance, it appears as if label "508" is mapped to "401" and label switching continues. When examining the label stack further, we see another process happening under the covers.

```
R5#Sh ip cef vrf TST 1.1.1.1 detail

1.1.1.1/32, version 21, epoch 0, cached adjacency to Serial1/0.54
0 packets, 0 bytes
  tag information set
    local tag: 508
    fast tag rewrite with Se1/0.54, point2point, tags imposed: {401 303}
  via 3.3.3.3, 0 dependencies, recursive
    next hop 10.1.45.4, Serial1/0.54 via 3.3.3.3/32
    valid cached adjacency
    tag rewrite with Se1/0.54, point2point, tags imposed: {401 303}
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The "401" label represents the top most label that will be forwarded to R4 to be label switched. The 303 (inner) label is the VPNV4 label that MP-BGP created on R3 and was signaled to R5. As you can see the "508" label is mapped to "401" in the VRF for forwarding purposes.

On R4

```
R4#Sh mpls forwarding-table | i 3.3.3.3

401     Pop tag     3.3.3.3/32         14059       Se0/0.43    point2point
```

The "pop tag" instruction will allow the inner label of "303" to be exposed when arriving at R3.

On R3

```
R3#Sh ip bgp vpnv4 all labels | i 1.1.1.1

   1.1.1.1/32       10.1.23.2       303/nolabel

R3#Sh mpls forwarding-table vrf TST | i 1.1.1.1

303     200         1.1.1.1/32[V]     0           Se1/0.32    point2point
```

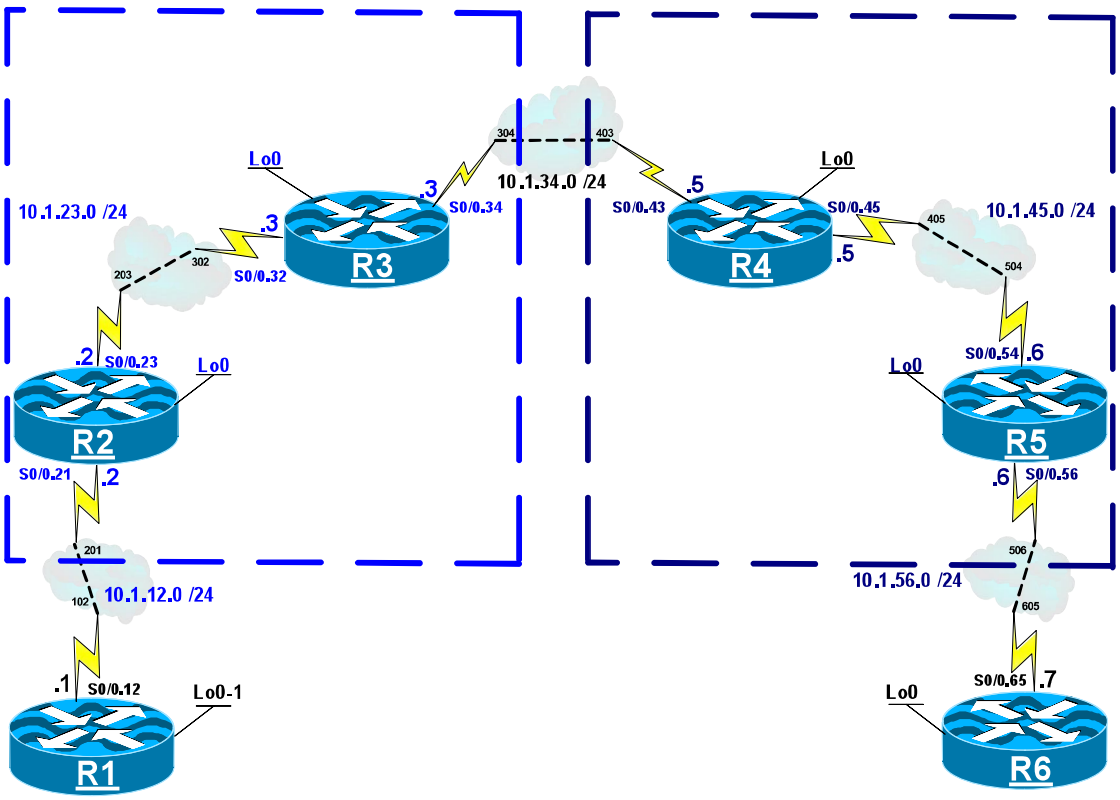
The "303/no label" displays the generation (control plane info) of this label by LDP and it is passed on to the MP-BGP process for delivery. This entry wants to strip the "303" label off. We can clearly see that after MPLS was configured with the CE, the action is now a label swap with "200" instead of being removed.

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 2
Inter-AS VPN



IP Addressing:

Router	Interface	IP address
R1	Lo0 Lo1 S0/0.12	1.1.1.1 /32 11.11.11.1/ 24 10.1.12.1 /24
R2	Lo0 S0/0.21 S0/0.23	2.2.2.2 /32 10.1.12.2 /24 10.1.23.2/ 24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

R3	Lo0 S0/0.32 S0/0.34	3.3.3.3 /32 10.1.23.3/ 24 10.1.34.3/ 24
R4	Lo0 S0/0.43 S0/0.45	4.4.4.4 /32 10.1.34.4/ 24 10.1.45.4 /24
R5	Lo0 S0/0.54 S0/0.56	5.5.5.5 /32 10.1.45.5/ 24 10.1.56.5/ 24
R6	Lo0 S0/0.65	6.6.6.6 /32 10.1.56.6/ /24

Task 1

Configure OSPF on all core routers. DO NOT allow OSPF connectivity to be possible between the PE and CE Routers R1 and R6.

On R2

```
R2(config)#router ospf 1
R2(config-router)#net 10.1.23.2 0.0.0.0 area 0
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#pass lo0
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.34.3 0.0.0.0 area 0
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#pass lo0
```

On R4

```
R4(config)#router ospf 1
R4(config-router)#netw 10.1.34.4 0.0.0.0 area 0
R4(config-router)#netw 10.1.45.4 0.0.0.0 area 0
R4(config-router)#netw 4.4.4.4 0.0.0.0 area 0
R4(config-router)#pass lo0
```

On R5

```
R5(config)#router ospf 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-router)#netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 0
R5(config-router)#pass lo0
```

Verify the Configuration:

On R2

```
R2#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O       3.3.3.3 [110/65] via 10.1.23.3, 00:01:05, Serial0/0.23
O       4.4.4.4 [110/129] via 10.1.23.3, 00:01:05, Serial0/0.23
O       5.5.5.5 [110/65] via 10.1.23.3, 00:01:05, Serial0/0.23
O       10.1.45.0 [110/192] via 10.1.23.3, 00:01:05, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:01:05, Serial0/0.23
```

On R5

```
R5#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O       2.2.2.2 [110/193] via 10.1.45.4, 00:00:17, Serial0/0.54
O       3.3.3.3 [110/129] via 10.1.45.4, 00:00:17, Serial0/0.54
O       4.4.4.4 [110/65] via 10.1.45.4, 00:00:17, Serial0/0.54
O       10.1.23.0 [110/192] via 10.1.45.4, 00:00:17, Serial0/0.54
O       10.1.34.0 [110/128] via 10.1.45.4, 00:00:17, Serial0/0.54
```

Task 2

Configure BGP on all core routers. R2 and R3 are a part of AS 65000. R4 and R5 are in AS 65001. Configure both AS's to support MP-BGP VPNv4 routes within and between them. There should be no MPLS configuration between AS's.

Create the backbone connections:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config)#mpls label range 200 299
R2(config)#mpls label protocol ldp
R2(config)#mpls ldp router-id lo0

R2(config)#int s0/0.23
R2(config-subif)#mpls ip

R2(config)#router bgp 65000
R2(config-router)#NO bgp default ipv4-unicast
R2(config-router)#neighbor 3.3.3.3 remot 65000
R2(config-router)#neighbor 3.3.3.3 update-source loop 0

R2(config-router)#address vpnv4 unicast
R2(config-router-af)#neighbor 3.3.3.3 activate
R2(config-router-af)#neighbor 3.3.3.3 send-community both
R2(config-router-af)#neighbor 3.3.3.3 next-hop-self
```

The "neighbor 3.3.3.3 send-community both" command is used for simplicity. We could have used "send-community extended" as the option if standard communities were not being used.

On R3

```
R3(config)#mpls label range 300 399
R3(config)#mpls label protocol ldp
R3(config)#mpls ldp router-id lo0

R3(config)#int s0/0.32
R3(config-subif)#mpls ip

R3(config)#ip bgp-community new-format
R3(config)#router bgp 65000
R3(config-router)#NO bgp default ipv4-uni
R3(config-router)#neighbor 2.2.2.2 remot 65000
R3(config-router)#neighbor 2.2.2.2 update loop 0
R3(config-router)#neighbor 10.1.34.4 remot 65001

R3(config-router)#address vpnv4
R3(config-router-af)#neighbor 2.2.2.2 activate
R3(config-router-af)#neighbor 2.2.2.2 send both
R3(config-router-af)#neighbor 2.2.2.2 next-hop-self
R3(config-router)#neighbor 10.1.34.4 activate
R3(config-router)#neighbor 10.1.34.4 send-community extended
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Here, the extended community is used toward the External BGP peer R4. Although we do send extended communities, the interface toward the EBGP peer is not configured to run MPLS LDP.

On R4

```
R4(config)#mpls label range 400 499
R4(config)#mpls label prot ldp
R4(config)#mpls ldp router-id lo0

R4(config)#int s0/0.45
R4(config-subif)#mpls ip

R4(config-subif)#ip bgp new-format
R4(config-subif)#router bgp 65001
R4(config-router)#NO bgp default ipv4-unicast
R4(config-router)#neighbor 5.5.5.5 remot 65001
R4(config-router)#neighbor 5.5.5.5 update lo0
R4(config-router)#neighbor 10.1.34.3 remot 65000

R4(config-router)#address vpnv4
R4(config-router-af)#neigh 5.5.5.5 activate
R4(config-router-af)#neigh 5.5.5.5 send both
R4(config-router-af)#neigh 5.5.5.5 next-hop-self
R4(config-router-af)#neighbor 10.1.34.3 activate
R4(config-router-af)#neighbor 10.1.34.3 send-community extended
```

On R5

```
R4(config)#mpls label range 500 599
R4(config)#mpls label prot ldp
R4(config)#mpls ldp router-id lo0

R4(config)#int s0/0.54
R4(config-subif)#mpls ip

R5(config-subif)#router bgp 65001
R5(config-router)#NO bgp default ipv4-uni
R5(config-router)#neighbor 4.4.4.4 remot 65001
R5(config-router)#neighbor 4.4.4.4 update lo0

R5(config-router)#address vpnv4
R5(config-router-af)#neighbor 4.4.4.4 activate
R5(config-router-af)#neighbor 4.4.4.4 send both
R5(config-router-af)#neighbor 4.4.4.4 next-hop-self
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Verify MP-BGP between Core Routers:

On R2

```
R2#Sh ip bgp vpnv4 all s | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4 65000      18      19        1    0    0 00:15:50        0
```

On R3

```
R3#Sh ip bgp vpnv4 all s | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4 65000      20      19        1    0    0 00:16:26        0
10.1.34.4       4 65001      10      11        1    0    0 00:07:17        0
```

On R4

```
R4#Sh ip bgp vpnv4 all s | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
5.5.5.5         4 65001      10      10        1    0    0 00:06:19        0
10.1.34.3       4 65000      11      10        1    0    0 00:07:37        0
```

On R5

```
R5#Sh ip bgp vpnv4 all s | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      10      10        1    0    0 00:06:49        0
```

Task 3

Create VRF's on the PE routers R2 and R5 with the following parameters:
VRF= TST, RD= 2:200 for R2 and 5:500 for R5, the RT for both Routers is 1:1000.

Create the VRF's:

On R2

```
R2(config)#ip vrf TST
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-vrf)#rd 2:200
R2(config-vrf)#route-target both 1:1000

R2(config-vrf)#int s0/0.21
R2(config-subif)#ip vrf forwarding TST
R2(config-subif)#ip addre 10.1.12.2 255.255.255.0
```

On R5

```
R5(config)#ip vrf TST
R5(config-vrf)#rd 5:500
R5(config-vrf)#route both 1:1000

R5(config-vrf)#int s0/0.56
R5(config-subif)#ip vrf forwarding TST
R5(config-subif)#ip add 10.1.56.5 255.255.255.0
```

Verify the configuration

On R2

```
R2#Sh ip vrf int

Interface          IP-Address      VRF                Protocol
Se0/0.21           10.1.12.2       TST                up

R2#Ping vrf TST 10.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 12/22/32 ms
```

On R5

```
R5#Sh ip vrf int

Interface                  IP-Address      VRF                Protocol
Se0/0.56                   10.1.56.5       TST                up

R5#Ping vrf TST 10.1.56.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.56.6, timeout is 2 seconds:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/31/52 ms
```

Task 4

Configure an OSPF customer context that uses process 20. Configure the lab in such a way that allows R1 and R6 to be able to ping each other's loopback interfaces without adding MPLS on the interface between providers.

On R5

```
R5(config)#router ospf 20 vrf TST
R5(config-router)#network 10.1.56.0 0.0.0.255 area 0
R5(config-router)#redistribute bgp 65001 subnets

R5(config-router)#router bgp 65001
R5(config-router)#address ipv4 vrf TST
R5(config-router-af)#redistribute ospf 20 match in ex 1 ex 2
```

On R2

```
R2(config)#router ospf 20 vrf TST
R2(config-router)#netw 10.1.12.0 0.0.0.255 area 0
R2(config-router)#redistribute bgp 65000 subnets

R2(config-router)#router bgp 65000
R2(config-router)#address ipv4 vrf TST
R2(config-router-af)#redistribute ospf 20 match internal external 1 ex 2
```

On R1

```
R1(config)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 11.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.12.0 0.0.0.255 area 0
```

On R6

```
R6(config)#router ospf 1
R6(config-router)#netw 10.1.56.0 0.0.0.255 area 0
R6(config-router)#netw 6.6.6.6 0.0.0.0 area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Verify and Test the Configuration:

On R6

```
R6#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The ping was not successful. Everything seems to have been configured correctly.

```
R6#Sh ip route | b Gateway

Gateway of last resort is not set

     6.0.0.0/32 is subnetted, 1 subnets
C       6.6.6.6 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.56.0 is directly connected, Serial0/0.65
```

The routing table of R6 backs up the output of the ping to ANY R1 prefix.

On R2

```
R2#Sh ip route vrf TST | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O       1.1.1.1 [110/65] via 10.1.12.1, 00:00:20, Serial0/0.21
O       11.11.11.1 [110/65] via 10.1.12.1, 00:00:20, Serial0/0.21
```

The routes are not seen from the other AS65000. A little more investigation will show where the problem resides.

```
R2#Sh ip bgp vpnv4 vrf TST | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 2:200 (default for vrf TST)
*> 1.1.1.1/32       10.1.12.1                65         32768 ?
*> 10.1.12.0/24     0.0.0.0                   0         32768 ?
*> 11.11.11.1/32    10.1.12.1                65         32768 ?
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The routes from R1 are in MP-BGP which rules out the route-targets as being the problem. The routes are not coming across the BGP session from R3.

On R5

```
R5#Sh ip bgp vpnv4 vrf TST

BGP table version is 5, local router ID is 5.5.5.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 5:500 (default for vrf TST)
*> 6.6.6.6/32       10.1.56.6                65         32768 ?
*> 10.1.56.0/24     0.0.0.0                   0         32768 ?
```

We confirm the BGP session seems to have an issue from R2 and R5 perspectives.
The routes appear as if they do not make it to R5.

```
R5#Sh ip bgp vpnv4 all summary | b Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      16      20        5    0    0 00:13:50        0
```

Routes even appear as if they do not exist in the update.

On R3

```
R3#Sh ip bgp vpnv4 all
R3#
```

The BGP VPNv4 routes are not showing up at all on R3. This is due to the fact that this router does not have the "TST" VRF configured and does not have the route-target "1:1000" anywhere. By default, BGP will filter ANY routes that do not contain the route-targets. We must turn off this filter to allow R3 to see the VPNv4 routes and propagate them to R4.

Remove the Route-Target Filter:

On R3

```
R3(config)#router bgp 65000
R3(config-router)#NO bgp default route-target filter

R3#Sh ip bgp vpnv4 all s | b Nei
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4 65000      31      22        4    0    0 00:19:17        3
10.1.34.4       4 65001      23      25        4    0    0 00:18:28        0
```

Performing the process one step at a time, we can see what happens when the filter is removed from R3 ONLY. It lifts the filter to display what R2 is sending. The BGP peer may need to be cleared using the " clear ip bgp * in" command to see the desired result.

```
R3#Sh ip bgp vpnv4 all | b Net
```

```
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 2:200
*>i1.1.1.1/32       2.2.2.2                  65    100      0 ?
*>i10.1.12.0/24     2.2.2.2                   0    100      0 ?
*>i11.11.11.1/32    2.2.2.2                  65    100      0 ?
```

The RD 2:200 is now visible. R3 can now see routes from R2. The routes from R4 cannot be seen yet due to the same filtering issue on R4.

On R4

```
R4#Sh ip bgp vpnv4 all s | b Nei
```

```
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
5.5.5.5         4 65001      32      29        6    0    0 00:23:14        0
10.1.34.3       4 65000      36      30        0    0    0 00:22:15        0
```

```
R4(config)#router bgp 65001
R4(config-router)#NO bgp default route-target filter

R4#Sh ip bgp vpnv4 all s | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
5.5.5.5         4 65001      32      29        6    0    0 00:23:38        2
10.1.34.3       4 65000      36      30        0    0    0 00:22:39        3
```

Now ALL of the VPNv4 routes can be seen.

On R6

```
R6#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O IA    1.1.1.1 [110/65] via 10.1.56.5, 00:02:07, Serial0/0.65
O IA    10.1.12.0 [110/65] via 10.1.56.5, 00:02:07, Serial0/0.65
O IA    11.11.11.1 [110/65] via 10.1.56.5, 00:02:07, Serial0/0.65
```

On R1

```
R1#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O IA    6.6.6.6 [110/65] via 10.1.12.2, 00:01:21, Serial0/0.12
O IA    10.1.56.0 [110/65] via 10.1.12.2, 00:01:21, Serial0/0.12


R1#Ping 6.6.6.6 source lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 104/140/168 ms
```

The forwarding is working properly for the MP-BGP routes.

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

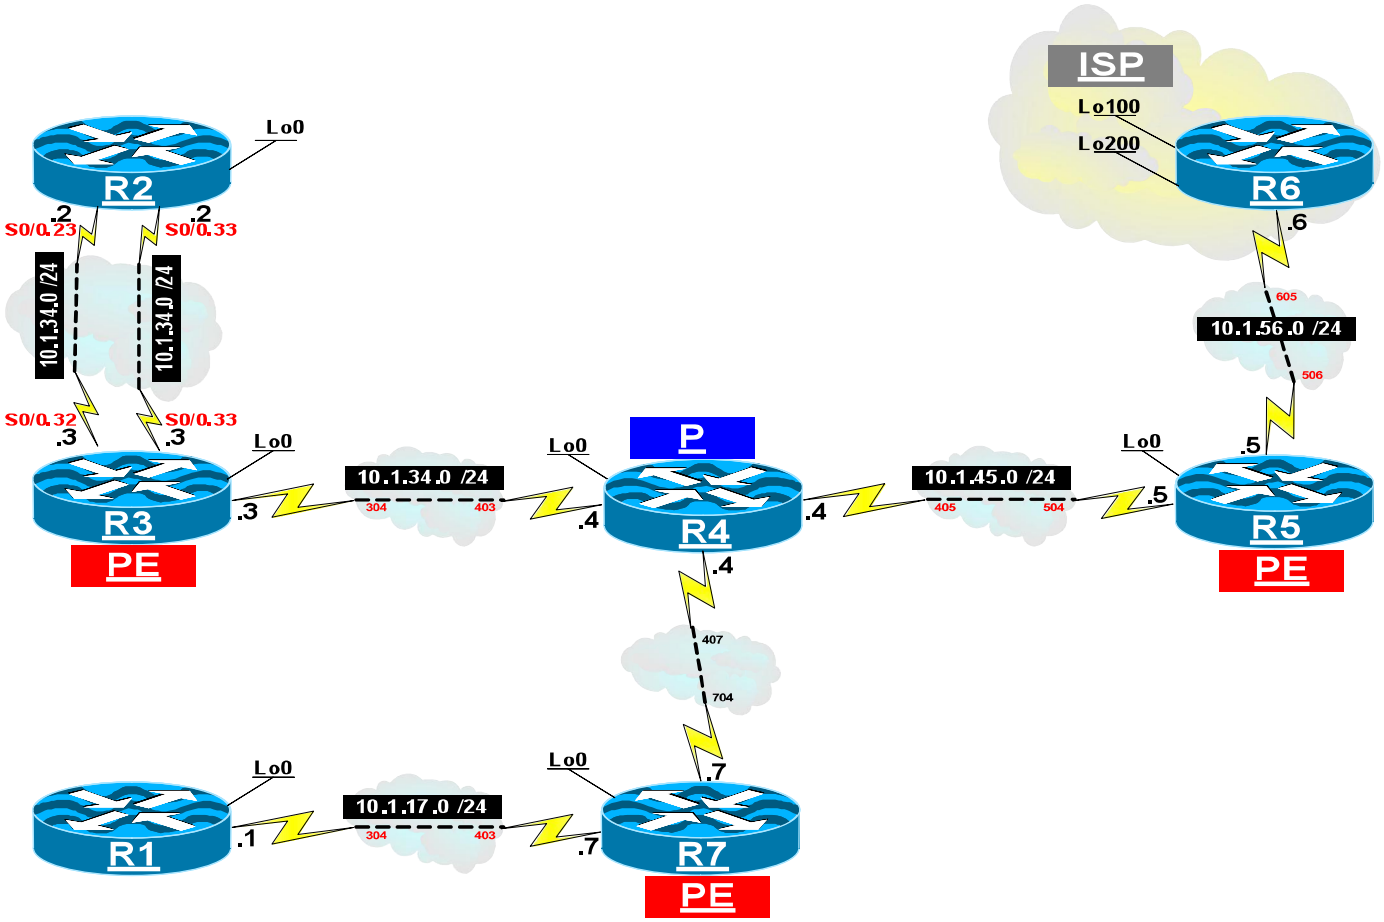
Paul Negron
CCIE #14856
SP

VPN Extranet & Internet Access

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

LAB 1-
Internet Access methods using separate VPN and dual PVC



Router	Interface	IP address
R1	Lo0 S0/0.17	1.1.1.1 /32 10.1.17.1 /24
R2	Lo0 F0/0 S0/0.23 S0/0.33	2.2.2.2/ 32 10.1.12.2 /24 10.1.23.2 /24 10.1.33.2/ 24
R3	Lo0 S0/0.32	3.3.3.3/ 32 10.1.23.3/ 24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

	S0/0.33 S0/0.34	10.1.33.3/ 24 10.1.34.3 /24
R4	Lo0 S0/0.43 S0/0.45 S0/0.47	4.4.4.4 /32 10.1.34.4/ 24 10.1.45.4/ 24 10.1.47.4 /24
R5	Lo0 S0/0.54 S0/0.56	5.5.5.5 /32 10.1.45.5 /24 10.1.56.5/ 24
R6	Lo100 Lo200 S0/0.65	100.100.100.100 /32 200.200.200.200 /32 10.1.56.6/ 24
R7	Lo0 S0/0.71 S0/0.74	7.7.7.7 /32 10.1.17.7 /24 10.1.47.7 /24

Lab Setup:

- Configure OSPF between Core routers R3, R4 and R5 using process "10"
- Configure BGP on R6 using an AS of 65066. An EBGP session shoud be preconfigured to peer with R5 10.1.56.5 address. Announce the Lo100 and Lo200 addresses into BGP.

Task 1

Configure Core Routers R3, R4, R5, and R7 to support MPLS VPN's using AS 65000. R5 should configure IPv4 peering to the physical interface of R6.

On R3

```
R3(config)#router ospf 10
R3(config-router)#network 0.0.0.0 0.0.0.0 area 0

R3(config)#mpls ldp router-id lo0
R3(config)#mpls label prot ldp

R3(config)#int s0/0.34
R3(config-subif)#mpls ip

R3(config)#router bgp 65000
R3(config-router)#NO bgp def ipv4
R3(config-router)#neighbor 5.5.5.5 remot 65000
R3(config-router)#neighbor 5.5.5.5 updat lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#neighbor 7.7.7.7 remot 65000
R3(config-router)#neighbor 7.7.7.7 updat lo0

R3(config-router)#address vpnv4
R3(config-router-af)#neighbor 5.5.5.5 act
R3(config-router-af)#neighbor 5.5.5.5 send ext
R3(config-router-af)#neighbor 5.5.5.5 next-hop-s
R3(config-router-af)#neighbor 7.7.7.7 act
R3(config-router-af)#neighbor 7.7.7.7 send ext
R3(config-router-af)#neighbor 7.7.7.7 next-hop-s
```

On R4

```
R4(config)#router ospf 10
R4(config-router)#network 0.0.0.0 0.0.0.0 area 0

R4(config)#mpls ldp router-id lo0
R4(config)#mpls label prot ldp

R4(config)#int s0/0.43
R4(config-subif)#mpls ip

R4(config-subif)#int s0/0.45
R4(config-subif)#mpls ip

R4(config-subif)#int s0/0.47
R4(config-subif)#mpls ip
```

On R5

```
R5(config)#router ospf 10
R5(config-router)#network 0.0.0.0 0.0.0.0 area 0

R5(config)#mpls ldp router-id lo0
R5(config)#mpls label prot ldp

R5(config)#int s0/0.54
R5(config-subif)#mpls ip

R5(config)#router bgp 65000
R5(config-router)#NO bgp def ipv4
R5(config-router)#neighbor 3.3.3.3 remot 65000
R5(config-router)#neighbor 3.3.3.3 updat lo0
R5(config-router)#neighbor 7.7.7.7 remot 65000
R5(config-router)#neighbor 7.7.7.7 updat lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-router)#address vpnv4
R5(config-router-af)#neighbor 3.3.3.3 act
R5(config-router-af)#neighbor 3.3.3.3 send ext
R5(config-router-af)#neighbor 3.3.3.3 next-hop-s
R5(config-router-af)#neighbor 7.7.7.7 act
R5(config-router-af)#neighbor 7.7.7.7 send ext
R5(config-router-af)#neighbor 7.7.7.7 next-hop-s
```

On R7

```
R7(config)#router ospf 10
R7(config-router)#network 0.0.0.0 0.0.0.0 area 0

R7(config)#mpls ldp router-id lo0
R7(config)#mpls label prot ldp

R7(config)#int s0/0.74
R7(config-subif)#mpls ip

R7(config)#router bgp 65000
R7(config-router)#NO bgp def ipv4
R7(config-router)#neighbor 5.5.5.5 remot 65000
R7(config-router)#neighbor 5.5.5.5 updat lo0
R7(config-router)#neighbor 3.3.3.3 remot 65000
R7(config-router)#neighbor 3.3.3.3 updat lo0

R7(config-router)#address vpnv4
R7(config-router-af)#neighbor 5.5.5.5 act
R7(config-router-af)#neighbor 5.5.5.5 send ext
R7(config-router-af)#neighbor 5.5.5.5 next-hop-s
R7(config-router-af)#neighbor 3.3.3.3 act
R7(config-router-af)#neighbor 3.3.3.3 send ext
R7(config-router-af)#neighbor 3.3.3.3 next-hop-s
```

On R5

```
R5(config)#router bgp 65000
R5(config-router)#neighbor 10.1.56.6 remot 65066

R5(config-router)#address ipv4
R5(config-router-af)#neighbor 10.1.56.6 act
R5(config-router-af)#neighbor 3.3.3.3 act

R5#Sh ip bgp s | b Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
3.3.3.3          4 65000      34       27         0    0    0 00:01:32 (NoNeg)
10.1.56.6        4 65066       9        8         3    0    0 00:04:06        2
```

The "NoNeg" State for BGP is shown when at least one Address-Family has been negotiated but one has not. In this case, it is the IPv4 Address-Family.

On R3

```
R3(config)#router bgp 65000
R3(config-router)#address ipv4
R3(config-router-af)#neighbor 5.5.5.5 act

R5#Sh ip bgp s | b Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.3         4   65000      41      33        3    0    0 00:00:18        0
10.1.56.6       4   65066      10       9        3    0    0 00:05:18        2
```

Task 2

Create a VRF called INTERNET on R3 and R7 and apply it to the directly connected interfaces toward R2 and R1. Use a RD of 3.3.3.3:100 and route-target of 65000:12 on R3 and a RD of 7.7.7.7:100 with the same RT configured on R3. R1 and R2 should be able to have ip reachability to ALL of the routes through this VPN. OSPF routing should be used for the VPN with a process of "12".

On R3

```
R3(config)#ip vrf INTERNET
R3(config-vrf)#rd 3.3.3.3:100
R3(config-vrf)#route 65000:12

R3(config-vrf)#int s0/0.33
R3(config-subif)#ip vrf forwarding INTERNET
R3(config-subif)#ip address 10.1.33.3 255.255.255.0

R3#Ping vrf INTERNET 10.1.33.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.33.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/16/44 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Testing basic connectivity in a VPN is always a good habit before any complicated configuration changes.

On R7

```
R7(config)#ip vrf INTERNET
R7(config-vrf)#rd 7.7.7.7:100
R7(config-vrf)#route 65000:12

R7(config-vrf)#int s0/0.71
R7(config-subif)#ip vrf forwarding INTERNET
R7(config-subif)#ip address 10.1.17.7 255.255.255.0

R7#Ping vrf INTERNET 10.1.17.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.17.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/26/56 ms
```

On R1

```
R1(config)#router ospf 12
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.17.0 0.0.0.255 area 0
```

On R2

```
R2(config)#router ospf 12
R2(config-router)#netw 2.2.2.2 0.0.0.0 area 0
R2(config-router)#netw 10.1.33.0 0.0.0.255 area 0
```

On R3

```
R3(config)#router ospf 12 vrf INTERNET
R3(config-router)#network 10.1.33.0 0.0.0.255 area 0
R3(config-router)#redistribute bgp 65000 subnets

R3(config-router)#router bgp 65000
R3(config-router)#address ipv4 vrf INTERNET
R3(config-router-af)#redistribute ospf 12 vrf INTERNET match in ex 1 ex 2
```

On R7

```
R7(config)#router ospf 12 vrf INTERNET
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R7(config-router)#network 10.1.17.0 0.0.0.255 area 0
R7(config-router)#redistribute bgp 65000 subnets

R7(config-router)#router bgp 65000
R7(config-router)#address ipv4 vrf INTERNET
R7(config-router-af)#redistribute ospf 12 vrf INTERNET match in ex 1 ex 2
```

On R1

```
R1#Sh ip route | in O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O IA    2.2.2.2 [110/129] via 10.1.17.7, 00:00:26, Serial0/0.17
O IA    10.1.33.0 [110/65] via 10.1.17.7, 00:00:26, Serial0/0.17


R1#Ping 2.2.2.2 source lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 48/133/228 ms
```

Task 3

Provide Internet Access for R1 in such a way that it uses R2. Static Routes are not permitted on R1 or R2. R2 should use BGP peering with an AS of 65002 to peer with R3.

On R2

```
R2(config)#router bgp 65002
R2(config-router)#neighbor 10.1.23.3 remot 65000
R2(config-router)#netw 1.1.1.1 mask 255.255.255.255
```

R2 advertises the OSPF route to the 1.1.1.1/32 prefix into BGP only.

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router bgp 65000
R3(config-router)#neighbor 10.1.23.2 remot 65002
R3(config-router)#address ipv4
R3(config-router-af)#neighbor 10.1.23.2 act

R3#Sh ip bgp | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.1/32       10.1.23.2              129             0 65002 i
*>i100.100.100.100/32
                    10.1.56.6                0    100      0 65066 i
*>i200.200.200.200/32
                    10.1.56.6                0    100      0 65066 i
```

R3 already has the routes that have come from R5 concerning the Internet Routing Table as well as the 1.1.1.1 prefix.

```
R2(config)#router ospf 12
R2(config-router)#default-information originate always
```

R2 sends the default route into the VPN to allow R1 to have a route to the internet. The "always" keyword is required due to the default route not being present.

On R1

```
R1#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O IA    2.2.2.2 [110/129] via 10.1.17.7, 00:17:30, Serial0/0.17
O IA    10.1.33.0 [110/65] via 10.1.17.7, 00:17:30, Serial0/0.17
```

R1 does not appear to have received the default route that was injected by R2 and therefore cannot reach the Internet. A little verification and Troubleshooting is in order to rectify the situation.

```
R3#Sh ip route vrf INTERNET | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O       2.2.2.2 [110/65] via 10.1.33.2, 00:00:54, Serial0/0.33
O*E2 0.0.0.0/0 [110/1] via 10.1.33.2, 00:00:54, Serial0/0.33
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The route appears to have reached the next hop router R3 in the VRF routing table.

On R3

```
R3#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 3.3.3.3:100 (default for vrf INTERNET)
*>i1.1.1.1/32       7.7.7.7                  65    100      0 ?
*> 2.2.2.2/32       10.1.33.2                65         32768 ?
*>i10.1.17.0/24     7.7.7.7                   0    100      0 ?
*> 10.1.33.0/24     0.0.0.0                   0         32768 ?
Route Distinguisher: 7.7.7.7:100
*>i1.1.1.1/32       7.7.7.7                  65    100      0 ?
*>i10.1.17.0/24     7.7.7.7                   0    100      0 ?
```

The default route was not redistributed into MP-BGP even though there is nothing wrong with the configuration on the surface.

```
R3(config)#router bgp 65000
R3(config-router)#address ipv4 vrf INTERNET
R3(config-router-af)#default-information originate
```

The process of redistributing a default route into MP-BGP is the problem and must be fixed by originating the default route into BGP.

```
R3#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 3.3.3.3:100 (default for vrf INTERNET)
*> 0.0.0.0          10.1.33.2                1         32768 ?
*>i1.1.1.1/32       7.7.7.7                  65    100      0 ?
*> 2.2.2.2/32       10.1.33.2                65         32768 ?
*>i10.1.17.0/24     7.7.7.7                   0    100      0 ?

*> 10.1.33.0/24     0.0.0.0                   0         32768 ?
Route Distinguisher: 7.7.7.7:100
*>i1.1.1.1/32       7.7.7.7                  65    100      0 ?
*>i10.1.17.0/24     7.7.7.7                   0    100      0 ?
```

The route is now passed to MP-BGP and should be propagated to the other PE Router R7.

On R1

```
R1#Sh ip route | i O
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O IA    2.2.2.2 [110/129] via 10.1.17.7, 00:19:55, Serial0/0.17
O IA    10.1.33.0 [110/65] via 10.1.17.7, 00:19:55, Serial0/0.17
```

The route is still not on R1.

```
R7#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 3.3.3.3:100
*>i0.0.0.0          3.3.3.3                  1    100      0 ?
*>i2.2.2.2/32       3.3.3.3                 65    100      0 ?
*>i10.1.33.0/24     3.3.3.3                  0    100      0 ?
Route Distinguisher: 7.7.7.7:100 (default for vrf INTERNET)
*>i0.0.0.0          3.3.3.3                  1    100      0 ?
*> 1.1.1.1/32       10.1.17.1               65          32768 ?
*>i2.2.2.2/32       3.3.3.3                 65    100      0 ?
*> 10.1.17.0/24     0.0.0.0                  0          32768 ?
*>i10.1.33.0/24     3.3.3.3                  0    100      0 ?
```

A closer look not only shows that the route made it over to R7 but it has also been properly maintained by MP-BGP. It is under both Route Distinguisher headings.

On R7

```
R7#Sh ip route vrf INTERNET | i B

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
B       2.2.2.2 [200/65] via 3.3.3.3, 00:21:06
B       10.1.33.0 [200/0] via 3.3.3.3, 00:21:20
B*   0.0.0.0/0 [200/1] via 3.3.3.3, 00:01:35
```

The route has also propagated into the VRF routing table, which proves that the import route-target statements are correct as well.

```
R7(config)#router ospf 12
R7(config-router)#default-information originate
```

A third default-information originate is required to finish the task. This time, it is performed under the OSPF context for the VRF.

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O IA     2.2.2.2 [110/129] via 10.1.17.7, 00:22:10, Serial0/0.17
O IA     10.1.33.0 [110/65] via 10.1.17.7, 00:22:10, Serial0/0.17
O*E2 0.0.0.0/0 [110/1] via 10.1.17.7, 00:00:24, Serial0/0.17

R1#Ping 200.200.200.200 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.200.200.200, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 240/298/360 ms
```

Task 4

Replace the route-targets applied to the INTERNET VRF on R7 with 65000:11. Create and apply the INTERNET VPN on R5 with a RD of 5.5.5.5:100 and a route-target of 555:555. Allow Internet Access from R5 to R1 and R2 by injecting a default route instead of both Internet routes currently viewed on R2. Pings should be successful from both VPN's to both Internet routes.

On R2

```
R2(config)#int s0/0.23
R2(config-subif)#shut
R2(config-subif)#router ospf 12

R2(config-router)#NO default-information originate
```

On R7

```
R7(config)#router ospf 12 vrf INTERNET
R7(config-router)#NO default-information originate

R7(config)#ip vrf INTERNET
R7(config-vrf)#NO route-target both 65000:12
R7(config-vrf)#route-target both 65000:11
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R5

```
R5(config)#ip vrf INTERNET
R5(config-vrf)#rd 5.5.5.5:100
R5(config-vrf)#route-target both 555:555

R5(config-vrf)#int s0/0.56
R5(config-subif)#ip vrf forwarding INTERNET
R5(config-subif)#ip address 10.1.56.5 255.255.255.0

R5#Ping vrf INTERNET 10.1.56.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.56.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/42/76 ms

R5(config)#router bgp 65000
R5(config-router)#address-family ipv4
R5(config-router-af)#NO neighbor 10.1.56.6
R5(config-router-af)#address-family ipv4 vrf INTERNET
R5(config-router-af)#neighbor 10.1.56.6 remot 65066
```

On R3 and R7

```
RX(config)#ip vrf INTERNET
RX(config-vrf)#route-target both 555:555
```

This is the new route-target that the ALL customers would participate in to receive Internet Routes and routes between them.

On R2

```
R2#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O E2    200.200.200.200 [110/1] via 10.1.33.3, 00:00:59, Serial0/0.33
O IA    1.1.1.1 [110/129] via 10.1.33.3, 00:04:46, Serial0/0.33
O E2    100.100.100.100 [110/1] via 10.1.33.3, 00:00:59, Serial0/0.33
O IA    10.1.17.0 [110/65] via 10.1.33.3, 00:04:46, Serial0/0.33
```

Although the 2 Internet routes are visible on R2, this will not be an acceptable site in a VPN. The

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

thought of potentially carrying all 300,000+ routes in the Internet will not scale in the VRF.

On R5

```
R5(config)#ip route vrf INTERNET 0.0.0.0 0.0.0.0 10.1.56.6

R5(config)#router bgp 65000
R5(config-router)#address ipv4 vrf INTERNET
R5(config-router-af)#netw 0.0.0.0
```

On R6

```
R6#Sh ip route | i 0.0.0.0

Gateway of last resort is 10.1.56.5 to network 0.0.0.0
     100.0.0.0/32 is subnetted, 1 subnets
     10.0.0.0/24 is subnetted, 3 subnets
B*   0.0.0.0/0 [20/0] via 10.1.56.5, 00:01:37
```

A default route must be injected into the VPN from R5 since the Task specifically stated not to have R2 see the Internet routes. R6 gets this route, which is not a very good side effect.

```
R5(config)#ip prefix-list NO-DEFAULT deny 0.0.0.0/0
R5(config)#ip prefix-list NO-DEFAULT permit 0.0.0.0/0 le 32

R5(config)#ip prefix-list DEFAULT-ONLY permit 0.0.0.0/0
```

A prefix-list is created to filter the default route from going to R6. The other prefix-list that allows the default only to come from R6 is not needed but is good practice to not allow a mistake by the provider to take place and introduce the entire BGP table from the Internet by mistake.

```
R5(config)#router bgp 65000
R5(config-router)#address-family ipv4 vrf INTERNET
R5(config-router-af)#neighbor 10.1.56.6 prefix-list DEFAULT-ONLY in
R5(config-router-af)#neighbor 10.1.56.6 prefix-list NO-DEFAULT out

R5#Clear ip bgp * in
R5#Clear ip bgp * out
```

On R6

```
R6#Sh ip route 0.0.0.0
% Network not in table
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 3.3.3.3:100 (default for vrf INTERNET)
*>i0.0.0.0          5.5.5.5                  0    100      0 i
*>i1.1.1.1/32       7.7.7.7                 65    100      0 ?
*> 2.2.2.2/32       10.1.33.2               65         32768 ?
*>i10.1.17.0/24     7.7.7.7                  0    100      0 ?
*> 10.1.33.0/24     0.0.0.0                  0         32768 ?
Route Distinguisher: 5.5.5.5:100
*>i0.0.0.0          5.5.5.5                  0    100      0 i
Route Distinguisher: 7.7.7.7:100
*>i1.1.1.1/32       7.7.7.7                 65    100      0 ?
*>i10.1.17.0/24     7.7.7.7                  0    100      0 ?
```

On R3 and R7

```
RX(config-vrf)#router ospf 12
RX(config-router)#default-info orig
```

On R2

```
R2#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O IA    1.1.1.1 [110/129] via 10.1.33.3, 02:28:41, Serial0/0.33
O IA    10.1.17.0 [110/65] via 10.1.33.3, 02:28:41, Serial0/0.33
O*E2 0.0.0.0/0 [110/1] via 10.1.33.3, 00:01:40, Serial0/0.33

R1#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area

       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O IA    2.2.2.2 [110/129] via 10.1.17.7, 02:30:45, Serial0/0.17
O IA    10.1.33.0 [110/65] via 10.1.17.7, 02:30:45, Serial0/0.17
O*E2 0.0.0.0/0 [110/1] via 10.1.17.7, 00:01:32, Serial0/0.17
```

The default route is has been sent to both CE routers as well as the VPN routes.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Ping 100.100.100.100 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.100.100, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 64/108/148 ms
```

On R2

```
R2#Ping 100.100.100.100 so lo0

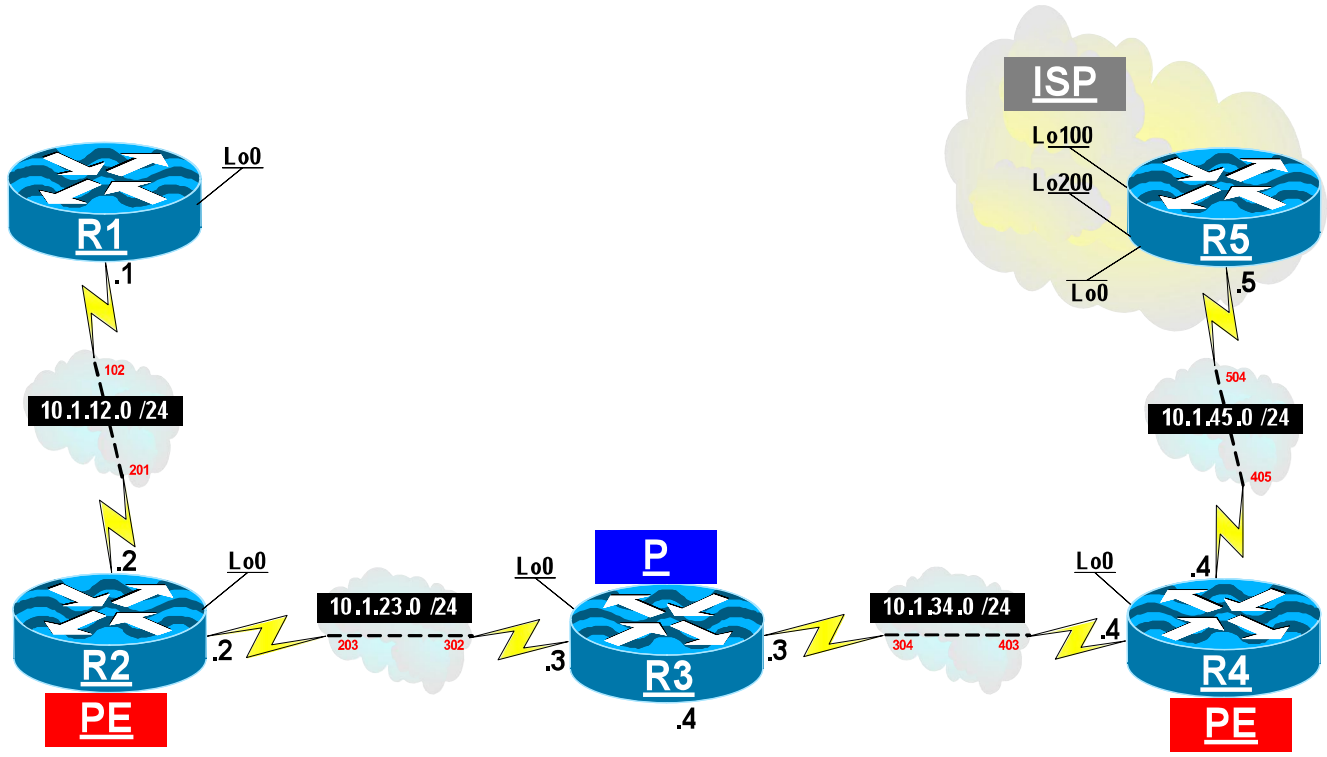
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.100.100.100, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 92/126/160 ms
```

Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

LAB 2-
Static Internet Access using Global Next Hop



IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12	1.1.1.1 /32 10.1.12.1 /24
R2	Lo0 S0/0.21 S0/0.23	2.2.2.2/ 32 10.1.12.2 /24 10.1.23.2 /24
R3	Lo0 S0/0.32 S0/0.34	3.3.3.3/ 32 10.1.23.3/ 24 10.1.34.3 /24
R4	Lo0 S0/0.43 S0/0.45	4.4.4.4 /32 10.1.34.4/ 24 10.1.45.4/ 24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

R5	Lo0 Lo100 Lo200 S0/0.54	5.5.5.5 /32 100.100.100.100 /32 200.200.200.200 /32 10.1.45.5 /24

Lab Setup:

- Configure OSPF between Core routers R2, R3 and R4 using process "1"
- Configure BGP on R5 using an AS of 65005. An EBGP session shoud be preconfigured to peer with R4 10.1.45.4 address. Announce the Lo100 and Lo200 addresses into BGP.

Task 1

Configure Core Routers R2, R3, and R4 to support MPLS VPN's. using AS 65234. R5 should configure IPv4 peering to the physical interface of R6.

On R2

```
R2(config)#mpls label prot ldp
R2(config)#mpls ldp router-id lo0

R2(config)#int s0/0.23
R2(config-subif)#mpls ip
```

On R3

```
R3(config)#mpls label prot ldp
R3(config)#mpls ldp router-id lo0

R3(config)#int s0/0.32
R3(config-subif)#mpls ip

R3(config)#int s0/0.34
R3(config-subif)#mpls ip
```

On R4

```
R4(config)#mpls label prot ldp
R4(config)#mpls ldp router-id lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#int s0/0.43
R4(config-subif)#mpls ip

R4(config)#router bgp 65234
R4(config-router)#NO bgp def ipv4
R4(config-router)#neighbor 10.1.45.5 remot 65005

R4(config-router)#address-family ipv4
R4(config-router-af)#neighbor 10.1.45.5 act

R4#Sh ip bgp vpnv4 all s | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4 65234       4       4        1    0    0 00:00:12        0

R4#Sh ip bgp s | b Nei

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.45.5       4 65005       6       5        3    0    0 00:01:13        2
```

Task 2

Create a VRF called INTERNET on R2 and apply it to the directly connected interface toward R1. Use an RD of 2.2.2.2:100 and route-target of 65234:11.

On R2

```
R2(config)#ip vrf INTERNET
R2(config-vrf)#rd 2.2.2.2:100
R2(config-vrf)#route 65234:11

R2(config-vrf)#int s0/0.21
R2(config-subif)#ip vrf forwarding INTERNET
R2(config-subif)#ip address 10.1.12.2 255.255.255.0

R2#Ping vrf INTERNET 10.1.12.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/17/28 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Provide Internet Access for R1's loopback 0 interface without using a dynamic routing protocol between the customer and the R2 PE Router.

On R1

```
R1(config)#ip route 0.0.0.0 0.0.0.0 10.1.12.2
```

The default route is configured on the CE since an IGP is not being used.

On R2

```
R2(config)#ip route vrf INTERNET 0.0.0.0 0.0.0.0 4.4.4.4 global
```

The default route is also configured on the PE to point to the Global Routing Table since R4 is not participating in the VPN in ANY capacity. This allows any packet coming from the CE to exit the VPN and reach a next hop in the Gloal routing Table for reachability to the Internet.

```
R2(config)#ip route vrf INTERNET 1.0.0.0 255.0.0.0 10.1.12.1
```

The loopback on R1 was specified to have connectivity to the Internet and is announced in the VRF to allow the PE router to reach this prefix.

```
R2(config)#ip route 1.0.0.0 255.0.0.0 s0/0.21 10.1.12.1
```

The same exact prefix is required to be configured in the global routing table to allow R4 to see this prefix when it is redistributed.

```
R2#Sh ip route vrf INTERNET | b Gateway
Gateway of last resort is 4.4.4.4 to network 0.0.0.0

S     1.0.0.0/8 [1/0] via 10.1.12.1
      10.0.0.0/24 is subnetted, 1 subnets
C        10.1.12.0 is directly connected, Serial0/0.21
S*    0.0.0.0/0 [1/0] via 4.4.4.4
```

On R4

```
R4#Sh ip route | b Gateway
Gateway of last resort is not set

      200.200.200.0/32 is subnetted, 1 subnets
B        200.200.200.200 [20/0] via 10.1.45.5, 00:10:14
      2.0.0.0/32 is subnetted, 1 subnets
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       2.2.2.2 [110/129] via 10.1.34.3, 00:08:54, Serial0/0.43
     100.0.0.0/32 is subnetted, 1 subnets
B       100.100.100.100 [20/0] via 10.1.45.5, 00:10:14
     3.0.0.0/32 is subnetted, 1 subnets
O       3.3.3.3 [110/65] via 10.1.34.3, 00:08:54, Serial0/0.43
     4.0.0.0/32 is subnetted, 1 subnets
C       4.4.4.4 is directly connected, Loopback0
     10.0.0.0/24 is subnetted, 3 subnets
O       10.1.23.0 [110/128] via 10.1.34.3, 00:08:54, Serial0/0.43
C       10.1.45.0 is directly connected, Serial0/0.45
C       10.1.34.0 is directly connected, Serial0/0.43
```

R4 does not see the route yet.

On R2

```
R2(config)#router ospf 1
R2(config-router)#redistribute static subnets

R4#Sh ip route | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route
O E2 1.0.0.0/8 [110/20] via 10.1.34.3, 00:00:13, Serial0/0.43
O       2.2.2.2 [110/129] via 10.1.34.3, 00:00:13, Serial0/0.43
O       3.3.3.3 [110/65] via 10.1.34.3, 00:00:13, Serial0/0.43
O       10.1.23.0 [110/128] via 10.1.34.3, 00:00:13, Serial0/0.43
```

The route is redistributed into the Global Instance of OSPF since redistributing this route into BGP will do us no good.

On R4

```
R4#Sh ip bgp | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
*> 100.100.100.100/32
                    10.1.45.5                0             0 65005 i
*> 200.200.200.200/32
                    10.1.45.5                0             0 65005 i
```

Although the route has reached R4, it still needs to be advertised into BGP to allow R6 to route back to the proper origin of the route, on R1.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#router bgp 65234
R4(config-router)#address ipv4
R4(config-router-af)#netw 1.0.0.0

R4#Sh ip bgp | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.34.3                20         32768 i
*> 100.100.100.100/32
                    10.1.45.5                 0             0 65005 i
*> 200.200.200.200/32
                    10.1.45.5                 0             0 65005 i
```

On R5

```
R5#Sh ip route | b Gateway
Gateway of last resort is not set

     200.200.200.0/32 is subnetted, 1 subnets
C       200.200.200.200 is directly connected, Loopback200
B    1.0.0.0/8 [20/20] via 10.1.45.4, 00:11:15
     100.0.0.0/32 is subnetted, 1 subnets
C       100.100.100.100 is directly connected, Loopback100
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/0.54
```

On R1

```
R1#Ping 200.200.200.200 source lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 200.200.200.200, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 44/68/84 ms
```

R1 can ping ANY Internet prefix now!!!

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

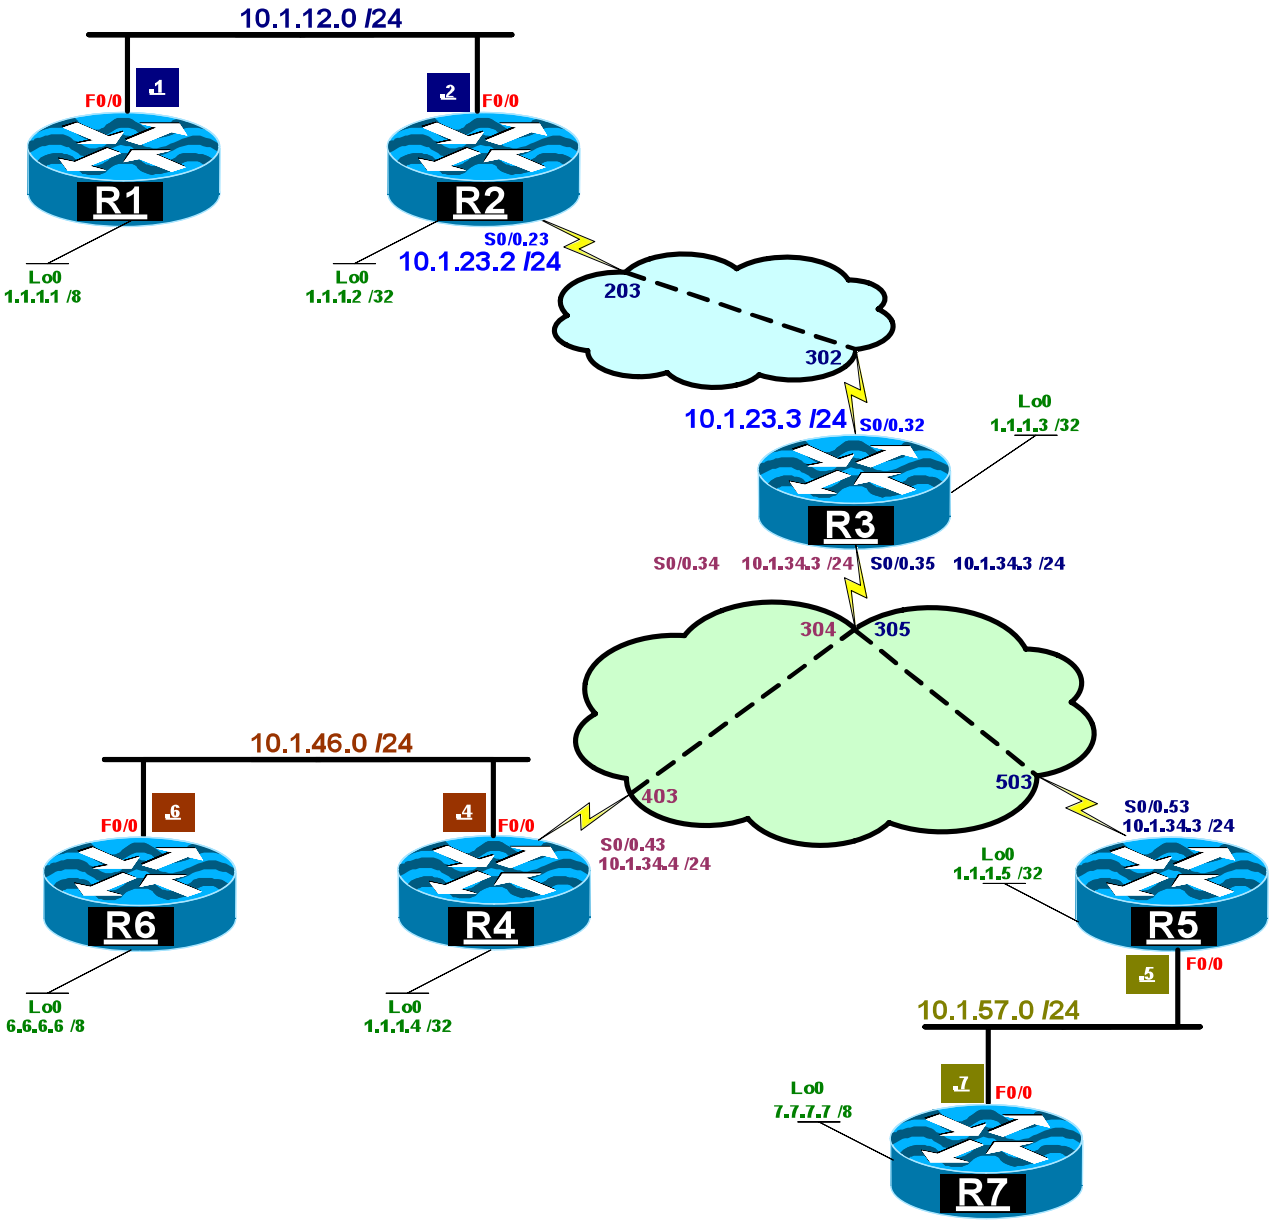
Paul Negron
CCIE #14856
SP

Multicast VPN (mVPN)

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1 – Multicast VPN (mVPN)



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 F0/0	1.1.1.1/ 8 10.1.12.1
R2	Lo0 F0/0 S0/0.23	1.1.1.2 /32 10.1.12.2/ 24 10.1.23.2 /24
R3	Lo0 S0/0.32 S0/0.35 S0/0.34	1.1.1.3/ 32 10.1.23.3 /24 10.1.35.3 /24 10.1.34.3/ 24
R4	Lo0 F0/0 S0/0.43	1.1.1.4 /32 10.1.46.4/ 24 10.1.34.4/ 24
R5	Lo0 F0/0 S0/0.53	1.1.1.5 /32 10.1.57.5/ 24 10.1.35.5/ 24
R6	Lo0 F0/0	6.6.6.6/8 10.1.46.6/ 24
R7	Lo0 F0/0	7.7.7.7/8 10.1.57.7/ 24

Lab Setup:

- The IP Addressing should reflect the table above

Task 1

Configure OSPF with a process of "20" on R2, R3, R4 and R5.

On R2

```
R2(config)#router ospf 20
R2(config-router)#network 0.0.0.0 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 20
R3(config-router)#network 0.0.0.0 0.0.0.0 area 0
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#router ospf 20
R4(config-router)#network 0.0.0.0 0.0.0.0 area 0
```

On R5

```
R5(config)#router ospf 20
R5(config-router)#network 0.0.0.0 0.0.0.0 area 0
```

Verifying the Configuration:

On R2

```
R2#Sh ip route ospf | i O

O       1.1.1.3 [110/65] via 10.1.23.3, 00:00:07, Serial0/0.23
O       1.1.1.5 [110/129] via 10.1.23.3, 00:00:07, Serial0/0.23
O       1.1.1.4 [110/129] via 10.1.23.3, 00:00:07, Serial0/0.23
O       10.1.46.0 [110/138] via 10.1.23.3, 00:00:07, Serial0/0.23
O       10.1.35.0 [110/128] via 10.1.23.3, 00:00:07, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:00:07, Serial0/0.23
O       10.1.57.0 [110/138] via 10.1.23.3, 00:00:07, Serial0/0.23
```

On R3

```
R3#Sh ip route ospf | i O

O       1.1.1.2/32 [110/65] via 10.1.23.2, 00:00:10, Serial0/0.32
O       1.1.1.5/32 [110/65] via 10.1.35.5, 00:00:10, Serial0/0.35
O       1.1.1.4/32 [110/65] via 10.1.34.4, 00:00:10, Serial0/0.34
O       10.1.12.0 [110/74] via 10.1.23.2, 00:00:10, Serial0/0.32
O       10.1.46.0 [110/74] via 10.1.34.4, 00:00:10, Serial0/0.34
O       10.1.57.0 [110/74] via 10.1.35.5, 00:00:10, Serial0/0.35
```

On R4

```
R4#Sh ip route ospf | i O

O       1.1.1.3 [110/65] via 10.1.34.3, 00:00:13, Serial0/0.43
O       1.1.1.2 [110/129] via 10.1.34.3, 00:00:13, Serial0/0.43
O       1.1.1.5 [110/129] via 10.1.34.3, 00:00:13, Serial0/0.43
O       10.1.12.0 [110/138] via 10.1.34.3, 00:00:13, Serial0/0.43
O       10.1.23.0 [110/128] via 10.1.34.3, 00:00:13, Serial0/0.43
O       10.1.35.0 [110/128] via 10.1.34.3, 00:00:13, Serial0/0.43
O       10.1.57.0 [110/138] via 10.1.34.3, 00:00:13, Serial0/0.43
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R5

```
R5#Sh ip route ospf | i O

O       1.1.1.3 [110/65] via 10.1.35.3, 00:00:14, Serial0/0.53
O       1.1.1.2 [110/129] via 10.1.35.3, 00:00:14, Serial0/0.53
O       1.1.1.4 [110/129] via 10.1.35.3, 00:00:14, Serial0/0.53
O       10.1.12.0 [110/138] via 10.1.35.3, 00:00:14, Serial0/0.53
O       10.1.23.0 [110/128] via 10.1.35.3, 00:00:14, Serial0/0.53
O       10.1.46.0 [110/138] via 10.1.35.3, 00:00:14, Serial0/0.53
O       10.1.34.0 [110/128] via 10.1.35.3, 00:00:14, Serial0/0.53
```

Task 2

Configure VPNv4 peering between PE routers R2,R4 and R5. MP-BGP should not allow version updates by default and should also use AS 1. Configure VRF's with a name of "TST" and the following attributes on the same routers:

Router	RD	RT
R2	1:10	1:1000
R4	1:20	1:1000
R5	1:30	1:1000

On R2

```
R2(config)#router bgp 1
R2(config-router)#NO bgp default ipv4
R2(config-router)#neighbor 1.1.1.4 remot 1
R2(config-router)#neighbor 1.1.1.4 update lo0
R2(config-router)#neighbor 1.1.1.5 remot 1
R2(config-router)#neighbor 1.1.1.5 update lo0

R2(config-router)#address vpnv4
R2(config-router-af)#neighbor 1.1.1.4 activate
R2(config-router-af)#neighbor 1.1.1.4 send both
R2(config-router-af)#neighbor 1.1.1.5 activate
R2(config-router-af)#neighbor 1.1.1.5 send both

R2(config-router)#ip vrf TST
R2(config-vrf)#rd 1:10
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-vrf)#route-target both 1:1000

R2(config-vrf)#int f0/0
R2(config-if)#ip vrf forwarding TST
R2(config-if)#ip add 10.1.12.2 255.255.255.0
```

On R4

```
R4(config)#router bgp 1
R4(config-router)#NO bgp default ipv4
R4(config-router)#neighbor 1.1.1.2 remot 1
R4(config-router)#neighbor 1.1.1.2 update lo0
R4(config-router)#neighbor 1.1.1.5 remot 1
R4(config-router)#neighbor 1.1.1.5 update lo0

R4(config-router)#address vpnv4
R4(config-router-af)#neighbor 1.1.1.2 activate
R4(config-router-af)#neighbor 1.1.1.2 send both
R4(config-router-af)#neighbor 1.1.1.5 activate
R4(config-router-af)#neighbor 1.1.1.5 send both
R4(config-router-af)#exit

R4(config-router)#ip vrf TST
R4(config-vrf)#rd 1:20
R4(config-vrf)#route-target both 1:1000

R4(config-vrf)#int f0/0
R4(config-if)#ip vrf forwarding TST
R4(config-if)#ip add 10.1.46.4 255.255.255.0
```

On R5

```
R5(config)#router bgp 1
R5(config-router)#NO bgp default ipv4
R5(config-router)#neighbor 1.1.1.2 remot 1
R5(config-router)#neighbor 1.1.1.2 update lo0
R5(config-router)#neighbor 1.1.1.4 remot 1
R5(config-router)#neighbor 1.1.1.4 update lo0

R5(config-router)#address vpnv4
R5(config-router-af)#neighbor 1.1.1.2 activate
R5(config-router-af)#neighbor 1.1.1.2 send both
R5(config-router-af)#neighbor 1.1.1.4 activate
R5(config-router-af)#neighbor 1.1.1.4 send both

R5(config)#ip vrf TST
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-vrf)#rd 1:30
R5(config-vrf)#route-target both 1:1000

R5(config-vrf)#int f0/0
R5(config-if)#ip vrf forwarding TST
R5(config-if)#ip add 10.1.57.5 255.255.255.0
```

Verify MP-BGP:

On R2

```
R2#Sh ip bgp vpnv4 all s | b Nei
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
1.1.1.4	4	1	7	8	1	0	0	00:02:30	0
1.1.1.5	4	1	6	6	1	0	0	00:01:17	0

On R4

```
R4#Sh ip bgp vpnv4 all s | b Nei
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
1.1.1.2	4	1	7	6	1	0	0	00:02:58	0
1.1.1.5	4	1	6	6	1	0	0	00:01:30	0

On R5

```
R5#Sh ip bgp vpnv4 all s | b Nei
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
1.1.1.2	4	1	6	6	1	0	0	00:01:49	0
1.1.1.4	4	1	6	6	1	0	0	00:01:35	0

Nothing fancy. Just a straight MPLS VPN configuration that will soon support mVPN as a service in addition to the MPLS VPN solution.

Task 3

Establish an EIGRP Neighbor adjacency with all CE Routers R1, R6 and R7. The Autonomous System should be the same used as BGP.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config-if)#router eigrp 1
R2(config-router)#NO auto

R2(config-router)#address-family ipv4 vrf TST
R2(config-router-af)#autonomous-system 1
R2(config-router-af)#NO auto
R2(config-router-af)#netw 10.0.0.0
R2(config-router-af)#redistribute bgp 1 metric 1 1 1 1 1

R2(config-router-af)#router bgp 1
R2(config-router)#address ipv4 vrf TST
R2(config-router-af)#redistribute eigrp 1
```

On R4

```
R4(config-if)#router eigrp 1
R4(config-router)#NO auto

R4(config-router)#address-family ipv4 vrf TST
R4(config-router-af)#NO auto
R4(config-router-af)#autonomous-system 1
R4(config-router-af)#network 10.0.0.0
R4(config-router-af)#redistribute bgp 1 metric 1 1 1 1 1

R4(config-router-af)#router bgp 1
R4(config-router)#address-family ipv4 vrf TST
R4(config-router-af)#redistribute eigrp 1
```

On R5

```
R5(config)#router eigrp 1
R5(config-router)#NO auto

R5(config-router)#address-family ipv4 vrf TST
R5(config-router-af)#NO auto
R5(config-router-af)#autonomous-system 1
R5(config-router-af)#network 10.0.0.0
R5(config-router-af)#redistribute bgp 1 metric 1 1 1 1 1

R5(config-router-af)#router bgp 1
R5(config-router)#address-family ipv4 vrf TST
R5(config-router-af)#redistribute eigrp 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1(config)#router eigrp 1
R1(config-router)#NO auto
R1(config-router)#netw 1.0.0.0
R1(config-router)#net 10.0.0.0
```

On R6

```
R6(config)#router eigrp 1
R6(config-router)#NO auto
R6(config-router)#netw 6.0.0.0
R6(config-router)#netw 10.0.0.0
```

On R7

```
R7(config)#router eigrp 1
R7(config-router)#NO auto
R7(config-router)#netw 7.0.0.0
R7(config-router)#net 10.0.0.0
```

Verify CE routing:

On R1

```
R1#Sh ip route eigrp | i D

D    6.0.0.0/8 [90/435200] via 10.1.12.2, 00:00:05, FastEthernet0/0
D    7.0.0.0/8 [90/435200] via 10.1.12.2, 00:00:20, FastEthernet0/0
D       10.1.46.0 [90/307200] via 10.1.12.2, 00:07:20, FastEthernet0/0
D       10.1.57.0 [90/307200] via 10.1.12.2, 00:06:20, FastEthernet0/0
```

On R6

```
R6#Sh ip route eigrp | i D

D    1.0.0.0/8 [90/435200] via 10.1.46.4, 00:00:01, FastEthernet0/0
D    7.0.0.0/8 [90/435200] via 10.1.46.4, 00:01:15, FastEthernet0/0
D       10.1.12.0 [90/307200] via 10.1.46.4, 00:08:12, FastEthernet0/0
D       10.1.57.0 [90/307200] via 10.1.46.4, 00:07:16, FastEthernet0/0
```

On R7

```
R7#Sh ip route eigrp | i D
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D    1.0.0.0/8 [90/435200] via 10.1.57.5, 00:00:18, FastEthernet0/0
D    6.0.0.0/8 [90/435200] via 10.1.57.5, 00:01:18, FastEthernet0/0
D       10.1.12.0 [90/307200] via 10.1.57.5, 00:07:59, FastEthernet0/0
D       10.1.46.0 [90/307200] via 10.1.57.5, 00:07:59, FastEthernet0/0
```

The ping would fail from the CE routers because LDP has not been configured between the PE and P routers.

Task 4

Configure mVPN with a default mdt of 239.1.1.1 and a data mdt of 224.1.2.3. All PE routers should allow trees to build to one another. R3 is not permitted to run any data mdt data but should allow the control mdt to pass through it.

On R2

```
R2(config)#ip vrf TST
R2(config-vrf)#mdt default 239.1.1.1
R2(config-vrf)#mdt data 224.1.2.0 0.0.0.255
R2(config-vrf)#exit
```

The wildcard mask is used to cover a range of groups that are permitted withing a VPN.

```
R2(config)#ip multicast-routing
R2(config)#ip multicast-routing vrf TST

R2(config)#int f0/0
R2(config-if)#ip pim sparse-dense-mode

R2(config-if)#int lo0
R2(config-if)#ip pim sparse-dense-mode

R2(config-if)#int s0/0.23
R2(config-subif)#ip pim sparse-dense-mode
R2(config-subif)#mpls ip
```

The VRF that was created earlier will have the control tree built with a group address of 239.1.1.1. This will allow for Multicast to create a GRE tunnel to ALL of the endpoints that share a common route-target in the VRF.

The multicast for the global configuration will ensure that R3 is able to participate with the control (default) tree setup. The MPLS encapsulation is only necessary for the return traffic from the

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

receiver to be successful when testing connectivity later in the lab with a ping.

On R4

```
R4(config)#ip vrf TST
R4(config-vrf)#mdt default 239.1.1.1
R4(config-vrf)#mdt data 224.1.2.0 0.0.0.255
R4(config-vrf)#exit

R4(config)#ip multicast-routing
R4(config)#ip multicast-routing vrf TST

R4(config)#int f0/0
R4(config-if)#ip pim sparse-dense-mode

R4(config-if)#int lo0
R4(config-if)#ip pim sparse-dense-mode

R4(config-if)#int s0/0.43
R4(config-subif)#ip pim sparse-dense-mode
R4(config-subif)#mpls ip
```

On R5

```
R5(config)#ip vrf TST
R5(config-vrf)#mdt default 239.1.1.1
R5(config-vrf)#mdt data 224.1.2.0 0.0.0.255
R5(config-vrf)#exit

R5(config)#ip multicast-routing
R5(config)#ip multicast-routing vrf TST

R5(config)#int f0/0
R5(config-if)#ip pim sparse-dense-mode

R5(config-if)#int lo0
R5(config-if)#ip pim sparse-dense-mode

R5(config-if)#int s0/0.53
R5(config-subif)#ip pim sparse-dense-mode
R5(config-subif)#mpls ip
```

On R3

```
R3(config)#ip multicast-routing
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#int s0/0.32
R3(config-subif)#ip pim sparse-dense-mode
R3(config-subif)#mpls ip

R3(config-subif)#int s0/0.34
R3(config-subif)#ip pim sparse-dense-mode
R3(config-subif)#mpls ip

R3(config-subif)#int s0/0.35
R3(config-subif)#ip pim sparse-dense-mode
R3(config-subif)#mpls ip
```

Verifying the default tree:

On R2

```
R2#Sh ip mroute

IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 00:04:55/stopped, RP 0.0.0.0, flags: DCZ
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0.23, Forward/Sparse-Dense, 00:00:54/00:00:00
    MVRF TST, Forward/Sparse-Dense, 00:04:55/00:00:00

(1.1.1.2, 239.1.1.1), 00:01:37/00:02:56, flags: T
  Incoming interface: Loopback0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0.23, Forward/Sparse-Dense, 00:00:54/00:00:00

(1.1.1.4, 239.1.1.1), 00:00:44/00:02:56, flags: TZ
  Incoming interface: Serial0/0.23, RPF nbr 10.1.23.3
  Outgoing interface list:
    MVRF TST, Forward/Sparse-Dense, 00:00:44/00:00:00

(1.1.1.5, 239.1.1.1), 00:00:36/00:02:56, flags: TZ
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Incoming interface: Serial0/0.23, RPF nbr 10.1.23.3
  Outgoing interface list:
    MVRF TST, Forward/Sparse-Dense, 00:00:42/00:00:00

(*, 224.0.1.40), 00:05:00/00:02:15, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0.23, Forward/Sparse-Dense, 00:01:00/00:00:00
    Loopback0, Forward/Sparse-Dense, 00:05:00/00:00:00
```

The SPT tree "T" as a flag is not new but the "Z" flag is unique to this configuration which denotes a Multicast Tunnel being created. The *"1.1.1.2, 239.1.1.1"* entry does not have this flag as it is originating locally on this router.

Task 5

All of the CE routers should now be configured to allow a group address that is joined on R1 f0/1 interface to reach a source from R7.

On R1

```
R1(config)#ip multicast-routing

R1(config)#int f0/0
R1(config-if)#ip pim sparse-dense-mode

R1(config-if)#int f0/1
R1(config-if)#ip pim sparse-dense-mode
R1(config-if)#ip igmp join-group 224.1.2.3

R1(config)#ip pim bidir-enable
```

On R7

```
R7(config)#ip multicast-routing

R7(config)#int lo0
R7(config-if)#ip pim sparse-dense-mode

R7(config-if)#int f0/0
R7(config-if)#ip pim sparse-dense-mode
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R7(config)#ip pim bidir-enable
```

On R6

```
R6(config)#ip multicast-routing

R6(config)#int f0/0
R6(config-if)#ip pim sparse-dense-mode

R6(config-if)#int lo0
R6(config-if)#ip pim sparse-dense-mode

R6(config)#ip pim bidir-enable
```

R6 is configured to show how the control tree and data tree act independently from each other.

Verifying the Configuration:

On R7

```
R7#Ping 224.1.2.3 source lo0

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 224.1.2.3, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.7

Reply to request 0 from 10.1.12.1, 84 ms
Reply to request 0 from 10.1.12.1, 84 ms
```

On R2

```
R2#Show ip mroute 239.1.1.1 summary | b \(

(*, 239.1.1.1), 00:18:25/stopped, RP 0.0.0.0, OIF count: 2, flags: DCZ
  (1.1.1.2, 239.1.1.1), 00:15:08/00:02:47, OIF count: 1, flags: T
  (1.1.1.4, 239.1.1.1), 00:14:14/00:02:47, OIF count: 1, flags: TZ
  (1.1.1.5, 239.1.1.1), 00:14:07/00:02:47, OIF count: 1, flags: TZ
```

The normal "sh ip mroute" command displays the default tree setup for the vrf endpoints.

On R5

```
R5#Sh ip mroute vrf TST 224.1.2.3 | b \(

(*, 224.1.2.3), 00:00:16/stopped, RP 0.0.0.0, flags: D
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:00:16/00:00:00
    Tunnel0, Forward/Sparse-Dense, 00:00:16/00:00:00

(7.7.7.7, 224.1.2.3), 00:00:16/00:02:44, flags: T
  Incoming interface: FastEthernet0/0, RPF nbr 10.1.57.7
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:16/00:00:00

(10.1.57.7, 224.1.2.3), 00:00:16/00:02:44, flags: T
  Incoming interface: FastEthernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:16/00:00:00

(*, 224.0.1.40), 00:08:54/00:02:34, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:04:28/00:00:00
    Tunnel0, Forward/Sparse-Dense, 00:08:55/00:00:00
```

The ping generated by R7 has created Multicast state to be created on R5 within the VRF. Notice how the Tunnel interface that was created as a GRE type now shows up in the OIL as a valid forwarding interface.

On R4

```
R4#Sh ip mroute vrf TST | b \(

(*, 224.1.2.3), 00:00:05/stopped, RP 0.0.0.0, flags: D
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:00:05/00:00:00
    Tunnel0, Forward/Sparse-Dense, 00:00:05/00:00:00

(7.7.7.7, 224.1.2.3), 00:00:05/00:02:57, flags: PT
  Incoming interface: Tunnel0, RPF nbr 1.1.1.5
  Outgoing interface list:
  FastEthernet0/0, Prune/Sparse-Dense, 00:00:05/00:02:54

(10.1.57.7, 224.1.2.3), 00:00:05/00:02:57, flags: PT
  Incoming interface: Tunnel0, RPF nbr 1.1.1.5
  Outgoing interface list:
    FastEthernet0/0, Prune/Sparse-Dense, 00:00:05/00:02:54

---text omitted-----
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4 the traffic is received on the tunnel but is eventually pruned toward R6 as there are no receivers joining this group at this time. If we would have used the "ip pim vrf spt-threshold infinity" command, R4 would never have switched to the shortest path and allowed the interface on F0/0 to foraward traffic at all. It would have been pruned immediately.

On R3

```
R3#Sh ip mroute summary | b \(

(*, 239.1.1.1), 00:03:36/stopped, RP 0.0.0.0, OIF count: 3, flags: D
  (1.1.1.2, 239.1.1.1), 00:03:36/00:02:53, OIF count: 2, flags: T
  (1.1.1.4, 239.1.1.1), 00:03:36/00:02:53, OIF count: 2, flags: T
  (1.1.1.5, 239.1.1.1), 00:03:36/00:02:53, OIF count: 2, flags: T
```

R3 only continues to see the control information. There is no VRF's terminating on this router. It is only a transit for the data tunnels. The interfaces in the OIL are forwarding Multiast Packets.

On R2

```
R2#Sh ip mroute vrf TST | b \(

(*, 224.1.2.3), 00:01:50/stopped, RP 0.0.0.0, flags: D
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:01:50/00:00:00
    Tunnel0, Forward/Sparse-Dense, 00:01:50/00:00:00

(7.7.7.7, 224.1.2.3), 00:01:50/00:01:12, flags: T
  Incoming interface: Tunnel0, RPF nbr 1.1.1.5
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:01:50/00:00:00

(10.1.57.7, 224.1.2.3), 00:01:50/00:01:12, flags: T
  Incoming interface: Tunnel0, RPF nbr 1.1.1.5
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:01:50/00:00:00
```

R2 accepts the traffic coming from the tunnel and floods it out of the F0/0 interface toward R1 since it does have a static join created on an interface.

Adding another receiver on R6:

On R6

```
R6(config)#interface lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R6(config-if)#ip igmp join-group 224.1.2.3
```

On R7

```
R7#Ping 224.1.2.3 source lo0

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 224.1.2.3, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.7

Reply to request 0 from 10.1.12.1, 44 ms
Reply to request 0 from 10.1.12.1, 128 ms
Reply to request 0 from 10.1.46.6, 112 ms
Reply to request 0 from 10.1.46.6, 108 ms
```

On R4

```
R4#Sh ip mroute vrf TST | b \(

(*, 224.1.2.3), 00:07:53/stopped, RP 0.0.0.0, flags: D
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:07:53/00:00:00
    Tunnel0, Forward/Sparse-Dense, 00:07:53/00:00:00

(7.7.7.7, 224.1.2.3), 00:02:20/00:02:16, flags: T
  Incoming interface: Tunnel0, RPF nbr 1.1.1.5
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:00:52/00:00:00

(10.1.57.7, 224.1.2.3), 00:02:20/00:02:16, flags: T
  Incoming interface: Tunnel0, RPF nbr 1.1.1.5
  Outgoing interface list:
    FastEthernet0/0, Forward/Sparse-Dense, 00:00:52/00:00:00
```

The interface on R4 has now trasitioned to a "forward" state.

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

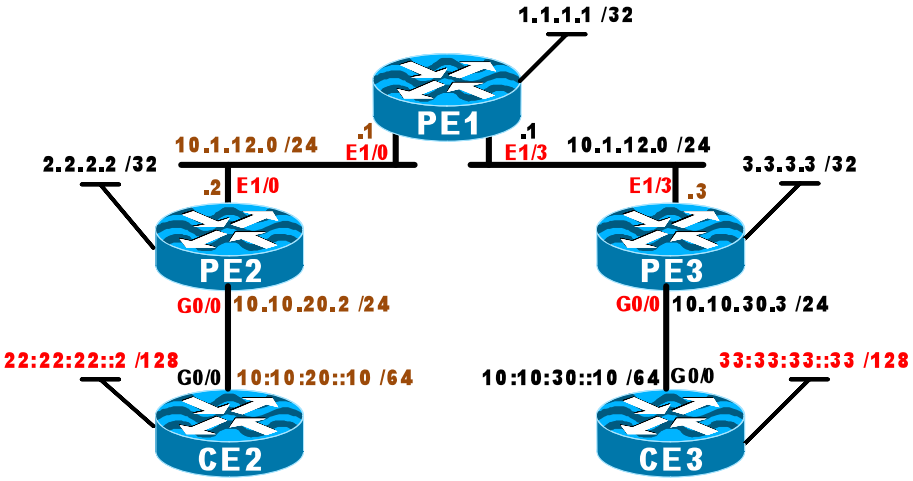
Paul Negron
CCIE #14856
SP

Layer 3 VPNs 6PE/6VPE

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

LAB 1- 6PE



Router	Interface	IPv4 address	IPv6 address	Link Local
PE1	Lo0 E1/0 E1/3	1.1.1.1 /32 10.1.12.1/ 24 10.1.13.1/ 24	---------- ---------- ----------	
PE2	Lo0 E1/0 G0/0	2.2.2.2 /32 10.1.12.2/ 24 10.10.20.2/ 24	---------- ---------- 10:10:20::2/64	 FE80::2
PE3	Lo0 E1/3 G0/0	3.3.3.3 /24 10.1.13.3/ 24 10.10.30.3/ 24	---------- ---------- 10:10:30::3/64	 FE80::3
CE2	Lo0 G0/0	---------- ----------	22:22:22::22/128 10:10:20::10/64	FE80::22 FE80::22
CE3	Lo0 G0/0	---------- ----------	33:33:33::33/128 10:10:30::10/64	FE80::33 FE80::33

Lab Setup:

- Configure the Ipv4 Addresses to the interfaces noted in the table above.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure the IPv6 Adresses on the appropriate interfaces. The routers should support ipv6 routing.

On PE2

```
PE2(config)#interface GigabitEthernet0/0
PE2(config-if)#ipv6 address FE80::2 link-local
PE2(config-if)#ipv6 address 10:10:20::2/64

PE2(config)#ipv6 unicast-routing
```

On PE3

```
PE3(config)#interface GigabitEthernet0/0
PE3(config-if)#ipv6 address FE80::3 link-local
PE3(config-if)#ipv6 address 10:10:30::3/64

PE3(config)#ipv6 unicast-routing
```

On CE2

```
CE2(config)#interface GigabitEthernet0/0
CE2(config-if)#ipv6 address FE80::22 link-local
CE2(config-if)#ipv6 address 10:10:20::10/64

CE2(config-if)#interface Lo0
CE2(config-if)#ipv6 address FE80::22 link-local
CE2(config-if)#ipv6 address 22:22:22::22/128

CE2(config)#ipv6 unicast-routing
```

On CE3

```
CE3(config)#interface GigabitEthernet0/0
CE3(config-if)#ipv6 address FE80::33 link-local
CE3(config-if)#ipv6 address 10:10:30::10/64

CE3(config-if)#interface Lo0
CE3(config-if)#ipv6 address FE80::33 link-local
CE3(config-if)#ipv6 address 33:33:33::33/128

CE3(config)#ipv6 unicast-routing
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The ipv6 unincast-routing command is not required to apply the address but will be required when performing static or dynamic routing.

Task 2

Configure ISIS as the routing protocol between core routers PE1,PE2 and PE3.

On PE2

```
PE2(config)#router isis TST
PE2(config-router)#net 49.0001.2222.0000.2222.00
PE2(config-router)#metric-style wide
PE2(config-router)#is-type level-2-only

PE2(config-router)#int e1/0
PE2(config-if)#ip router isis TST

PE2(config-if)#int lo0
PE2(config-if)#ip router isis TST
```

On PE1

```
PE1(config)#router isis TST
PE1(config-router)#net 49.0001.1111.0000.1111.00
PE1(config-router)#metric-style wide
PE1(config-router)#is-type level-2-only

PE1(config-if)#int lo0
PE1(config-if)#ip router isis TST
PE1(config-if)#int e1/0
PE1(config-if)#ip router isis TST
PE1(config-if)#int e1/3
PE1(config-if)#ip router isis TST
```

On PE3

```
PE3(config)#router isis TST
PE3(config-router)#net 49.0001.3333.0000.3333.00
PE3(config-router)#metric-style wide
PE3(config-router)#is-type level-2-only

PE3(config-router)#int lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
PE3(config-if)#ip router isis TST

PE3(config-if)#int e1/3
PE3(config-if)#ip router isis TST
```

Verifying Connectivity:

On PE3

```
PE3#Sh ip route isis | I L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2     1.1.1.1 [115/20] via 10.1.13.1, 00:00:30, Ethernet1/3
i L2     2.2.2.2 [115/30] via 10.1.13.1, 00:00:20, Ethernet1/3
i L2     10.1.12.0/24 [115/20] via 10.1.13.1, 00:00:30, Ethernet1/3
```

On PE1

```
PE1#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2     2.2.2.2 [115/20] via 10.1.12.2, 00:01:39, Ethernet1/0
i L2     3.3.3.3 [115/20] via 10.1.13.3, 00:00:52, Ethernet1/3
```

On PE2

```
PE2#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2     1.1.1.1 [115/20] via 10.1.12.1, 00:01:09, Ethernet1/0
i L2     3.3.3.3 [115/30] via 10.1.12.1, 00:00:19, Ethernet1/0
i L2     10.1.13.0/24 [115/20] via 10.1.12.1, 00:01:09, Ethernet1/0

PE2#Ping 3.3.3.3 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/48/56 ms
```

Task 3

Configure the backbone to support MPLS connectivity between PE2 and PE3 and BGP IPv6 peering.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

On PE2

```
PE2(config)#router bgp 123
PE2(config-router)#No bgp default ipv4
PE2(config-router)#neighbor 3.3.3.3 remot 123
PE2(config-router)#neighbor 3.3.3.3 update lo0

PE2(config-router)#address ipv6
PE2(config-router-af)#neighbor 3.3.3.3 act
```

On PE3

```
PE3(config)#router bgp 123
PE3(config-router)#No bgp default ipv4
PE3(config-router)#neighbor 2.2.2.2 remot 123
PE3(config-router)#neighbor 2.2.2.2 update lo0

PE3(config-router)#address ipv6
PE3(config-router-af)#neighbor 2.2.2.2 act
PE3(config-router-af)#end

PE3#Sh ip bgp ipv6 uni summary | b Nei

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4          123       4        4        1    0    0 00:00:22        0
```

Unlike a standard MPLS VPN backbone that requires VPNv4 peering, a 6 PE configuration does not require the backbone to support Route Distinguishers or VPNv4 peering to accomplish the tunneling of IPv6 traffic.

On PE1

```
PE1(config)#mpls label range 100 199
PE1(config)#mpls label prot ldp
PE1(config)#mpls ldp router-id lo0

PE1(config)#int e1/0
PE1(config-if)#mpls ip
PE1(config-if)#int e1/3
PE1(config-if)#mpls ip
```

On PE2

```
PE2(config)#mpls label range 200 299
PE2(config)#mpls label prot ldp
PE2(config)#mpls ldp router-id lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
PE2(config)#int e1/0
PE2(config-if)#mpls ip
```

On PE3

```
PE3(config)#mpls label range 300 399
PE3(config)#mpls label prot ldp
PE3(config)#mpls ldp router-id lo0

PE3(config)#int e1/3
PE3(config-if)#mpls ip

PE3(config-if)#end
```

LDP does not currently support the creation of labels based on IGP routes in the backbone using strict IPv6 protocols.

Task 4

Create a connection between CE2 and CE3 through the MPLS backbone. No GRE Tunnels are allowed.

On PE2

```
PE2(config)#router bgp 123
PE2(config-router)#address-family ipv6
PE2(config-router-af)#network 10:10:20::/64
```

On PE3

```
PE3(config)#router bgp 123
PE3(config-router)#address ipv6
PE3(config-router-af)#netw 10:10:30::/64
```

Verifying the Configuration

On PE2

```
PE2#Sh mpls forwarding-table

Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
200         Pop Label  1.1.1.1/32       0              Et1/0       10.1.12.1
201         101        3.3.3.3/32       0              Et1/0       10.1.12.1
202         Pop Label  10.1.13.0/24      0              Et1/0       10.1.12.1
```

The top most label that represents the path through the core is still implemented as an IPv4 tunnel using LDP.

```
PE2#Sh ip bgp ipv6 uni labels

   Network          Next Hop      In label/Out label
   10:10:20::/64    ::              nolabel/nolabel
   10:10:30::/64    ::FFFF:3.3.3.3  nolabel/nolabel

PE2#Sh ip bgp ipv6 uni neighbors | s Neighbor capabilities

  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv6 Unicast: advertised and received
    Multisession Capability: advertised and received
```

The prefix is learned through MP-BGP as an IPv6 route. The next hop is also seen as the IPv4 mapped address format of the loopback on the terminating router. The label however is not created due to a lack of BGP understanding a new capability.

```
PE2(config)#router bgp 123
PE2(config-router)#address-family ipv6
PE2(config-router-af)#neighbor 3.3.3.3 send-label

*Mar 25 12:19:49.735: %BGP-5-ADJCHANGE: neighbor 3.3.3.3 Down Capability changed
*Mar 25 12:19:49.739: %BGP_SESSION-5-ADJCHANGE: neighbor 3.3.3.3 IPv6 Unicast topology
base removed from session  Capability changed
```

The "send-label" command is configured to allow BGP to create and send a local label for the IPv6 prefix. This command only allows for a label to be advertised so it must be configured on both endpoints. The label will not be created until both peers have successfully negotiated the "ipv6 MPLS Label capability".

```
PE2#Sh ip bgp ipv6 uni neighbors | s Neighbor capabilities

  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv6 Unicast: advertised and received
    ipv6 MPLS Label capability: advertised
    Multisession Capability: advertised and received
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On PE3

```
PE3#Sh ip bgp ipv6 uni neighbors | s Neighbor capabilities

  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv6 Unicast: advertised and received
    ipv6 MPLS Label capability: received
    Multisession Capability: advertised and received

PE3#Sh ip bgp ipv6 uni labels

   Network          Next Hop      In label/Out label
   10:10:20::/64    ::FFFF:2.2.2.2  nolabel/nolabel
   10:10:30::/64    ::              nolabel/nolabel

PE3(config)#router bgp 123
PE3(config-router)#address-family ipv6
PE3(config-router-af)#neighbor 2.2.2.2 send-label

PE3#Sh ip bgp ipv6 uni neighbors | s Neighbor capabilities

  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv6 Unicast: advertised and received
    ipv6 MPLS Label capability: advertised and received
    Multisession Capability: advertised and received

PE3#Sh ip bgp ipv6 uni labels

   Network          Next Hop      In label/Out label
   10:10:20::/64    ::FFFF:2.2.2.2  nolabel/203
   10:10:30::/64    ::              303/nolabel
```

The inner label representing the path is now created and propagated.

On PE2

```
PE2#Sh ip bgp ipv6 uni labels

   Network          Next Hop      In label/Out label
   10:10:20::/64    ::              203/nolabel
   10:10:30::/64    ::FFFF:3.3.3.3  nolabel/303
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
PE2#Sh mpls forwarding-table labels 203

Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface

203        No Label   10:10:20::/64    0             aggregate
```

On CE2

```
CE2(config)#ipv6 route ::/0 g0/0 FE80::2
```

On CE3

```
CE3(config)#ipv6 route ::/0 g0/0 FE80::3

CE3#Ping 10:10:20::10

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10:10:20::10, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/47/104 ms
```

On PE3

```
PE3#Sh mpls forwarding-table 10:10:20::10/64 detail

Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
None       203        10:10:20::/64    0             Et1/3      10.1.13.1
     MAC/Encaps=14/22, MRU=1496, Label Stack{100 203}
     CA000791001FCA020791001F8847 00064000000CB000
     No output feature configured

CE3#Trace 10:10:20::10

Type escape sequence to abort.
Tracing the route to 10:10:20::10

  1 10:10:30::3 36 msec 4 msec 8 msec
  2 ::FFFF:10.1.13.1 [MPLS: Labels 100/203 Exp 0] 24 msec 16 msec 24 msec
  3 10:10:20::2 32 msec 8 msec 12 msec
  4 10:10:20::10 60 msec 8 msec 12 msec

PE3#Sh ip bgp ipv6 uni labels

   Network          Next Hop      In label/Out label
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   10:10:20::/64     ::FFFF:2.2.2.2  nolabel/203
   10:10:30::/64     ::              303/nolabel
```

Although the traceroute notes the IPv4 mapped address of the physical interface, the destination will see the interface of the PE router as the source. The Data Plane reveals tha label in the stack being applied with label "203" as the inner tag (ipv6) and "100" as the outer tag (ipv4).

```
PE3#Sh mpls forwarding-table

Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
300        Pop Label  1.1.1.1/32       0             Et1/3       10.1.13.1
301        100        2.2.2.2/32       0             Et1/3       10.1.13.1
302        Pop Label  10.1.12.0/24     0             Et1/3       10.1.13.1
303        No Label   10:10:30::/64    456           aggregate
```

On PE2

```
PE2#Sh ip bgp ipv6 uni labels

   Network          Next Hop      In label/Out label
   10:10:20::/64    ::              203/nolabel
   10:10:30::/64    ::FFFF:3.3.3.3  nolabel/303
```

On CE3

```
CE3#Ping 10:10:20::10 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10:10:20::10, timeout is 2 seconds:
Packet sent with a source address of 33:33:33::33
.....
Success rate is 0 percent (0/5)
```

There is no reachibility between the loopbacks of the CE's due to a lack of label creation for these prefixes.

Task 5

Allow for CE2 and CE3 to have reachability between the loopbacks using BGP as the Routing Protocol. CE2 should use AS "200" and CE3 should use AS "300".

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On CE2

```
CE2(config)#NO ipv6 route ::/0 g0/0 FE80::2

CE2(config)#router bgp 65200
CE2(config-router)#bgp router-id 22.22.22.22
CE2(config-router)#No bgp default ipv4
CE2(config-router)#neighbor 10:10:20::2 remot 123
CE2(config-router)#address ipv6
CE2(config-router-af)#neighbor 10:10:20::2 act
CE2(config-router-af)#netw 22:22:22::22/128
```

When using pure IPv6 BGP peering, a Router-ID is required. This address is configured as an IPv4 prefix.

On PE2

```
PE2(config)#router bgp 123
PE2(config-router)#neighbor 10:10:20::10 remot 65200
PE2(config-router)#address ipv6
PE2(config-router-af)#neighbor 10:10:20::10 act
```

On CE3

```
CE3(config)#NO ipv6 route ::/0 g0/0 FE80::3

CE3(config)#router bgp 65300
CE3(config-router)#bgp router-id 33.33.33.33
CE3(config-router)#No bgp default ipv4
CE3(config-router)#neighbor 10:10:30::3 remot 123
CE3(config-router)#address ipv6
CE3(config-router-af)#neighbor 10:10:30::3 act
CE3(config-router-af)#netw 33:33:33::33/128
```

On PE3

```
PE3(config)#router bgp 123
PE3(config-router)#neighbor 10:10:30::10 remot 65300
PE3(config-router)#address ipv6
PE3(config-router-af)#neighbor 10:10:30::10 act
```

On CE3

```
CE3#Sh ip bgp ipv6 uni | b Netw
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 10:10:20::/64    10:10:30::3                             0 123 i
*> 10:10:30::/64    10:10:30::3                0            0 123 i
*> 22:22:22::22/128 10:10:30::3                             0 123 65200 i
*> 33:33:33::33/128 ::                         0        32768 i
```

On PE3

```
PE3#Sh ip bgp ipv6 uni labels

   Network          Next Hop      In label/Out label
   10:10:20::/64    ::FFFF:2.2.2.2  nolabel/203
   10:10:30::/64    ::              303/nolabel
   22:22:22::22/128 ::FFFF:2.2.2.2  nolabel/204
   33:33:33::33/128 10:10:30::10    304/nolabel
```

On CE2

```
CE2#Sh ipv6 route bgp | i B

       B - BGP, R - RIP, I1 - ISIS L1, I2 - ISIS L2
B   10:10:30::/64 [20/0]
B   33:33:33::33/128 [20/0]

CE2#Ping 33:33:33::33 source lo0

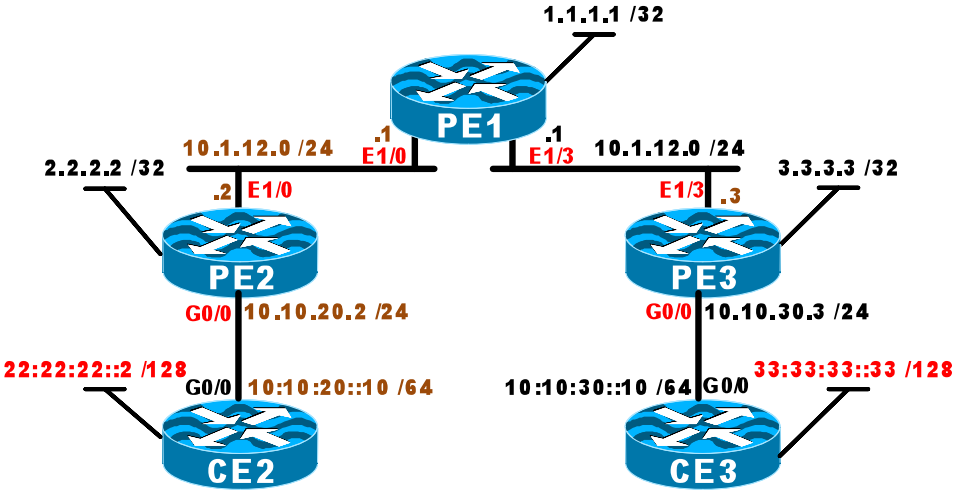
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 33:33:33::33, timeout is 2 seconds:
Packet sent with a source address of 22:22:22::22
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 20/34/40 ms
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

LAB 2- 6VPE



Router	Interface	IPv4 address	IPv6 address	Link Local
PE1	Lo0 E1/0 E1/3	1.1.1.1 /32 10.1.12.1/ 24 10.1.13.1/ 24	---------- ---------- ----------	
PE2	Lo0 E1/0 G0/0	2.2.2.2 /32 10.1.12.2/ 24 10.10.20.2/ 24	---------- ---------- 10:10:10::2/64	 FE80::2
PE3	Lo0 E1/3 G0/0	3.3.3.3 /24 10.1.13.3/ 24 10.10.30.3/ 24	---------- ---------- 10:10:30::3/64	 FE80::3
CE2	Lo0 G0/0	---------- ----------	22:22:22::22/128 10:10:20::10/64	FE80::22 FE80::22
CE3	Lo0 G0/0	---------- ----------	33:33:33::33/128 10:10:30::10/64	FE80::33 FE80::33

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab Setup:

- Configure the IPv4 Addresses to the interfaces noted in the table above.

Task 1

Configure the IPv6 Adresses on the appropriate interfaces. The routers should support ipv6 routing.

On PE2

```
PE2(config)#interface GigabitEthernet0/0
PE2(config-if)#ipv6 address FE80::2 link-local
PE2(config-if)#ipv6 address 10:10:20::2/64

PE2(config)#ipv6 unicast-routing
```

On PE3

```
PE3(config)#interface GigabitEthernet0/0
PE3(config-if)#ipv6 address FE80::3 link-local
PE3(config-if)#ipv6 address 10:10:30::3/64

PE3(config)#ipv6 unicast-routing
```

On CE2

```
CE2(config)#interface GigabitEthernet0/0
CE2(config-if)#ipv6 address FE80::22 link-local
CE2(config-if)#ipv6 address 10:10:20::10/64

CE2(config-if)#interface Lo0
CE2(config-if)#ipv6 address FE80::22 link-local
CE2(config-if)#ipv6 address 22:22:22::22/128

CE2(config)#ipv6 unicast-routing
```

On CE3

```
CE3(config)#interface GigabitEthernet0/0
CE3(config-if)#ipv6 address FE80::33 link-local
CE3(config-if)#ipv6 address 10:10:30::10/64

CE3(config-if)#interface Lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
CE3(config-if)#ipv6 address FE80::33 link-local
CE3(config-if)#ipv6 address 33:33:33::33/128

CE3(config)#ipv6 unicast-routing
```

Task 2

Configure ISIS as the routing protocol between core routers PE1,PE2 and PE3.

On PE2

```
PE2(config)#router isis TST
PE2(config-router)#net 49.0001.2222.0000.2222.00
PE2(config-router)#metric-style wide
PE2(config-router)#is-type level-2-only

PE2(config-router)#int e1/0
PE2(config-if)#ip router isis TST

PE2(config-if)#int lo0
PE2(config-if)#ip router isis TST
```

On PE1

```
PE1(config)#router isis TST
PE1(config-router)#net 49.0001.1111.0000.1111.00
PE1(config-router)#metric-style wide
PE1(config-router)#is-type level-2-only

PE1(config-if)#int lo0
PE1(config-if)#ip router isis TST
PE1(config-if)#int e1/0
PE1(config-if)#ip router isis TST
PE1(config-if)#int e1/3
PE1(config-if)#ip router isis TST
```

On PE3

```
PE3(config)#router isis TST
PE3(config-router)#net 49.0001.3333.0000.3333.00
PE3(config-router)#metric-style wide
PE3(config-router)#is-type level-2-only
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
PE3(config-router)#int lo0
PE3(config-if)#ip router isis TST

PE3(config-if)#int e1/3
PE3(config-if)#ip router isis TST
```

Verifying Connectivity:

On PE3

```
PE3#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      1.1.1.1 [115/20] via 10.1.13.1, 00:00:30, Ethernet1/3
i L2      2.2.2.2 [115/30] via 10.1.13.1, 00:00:20, Ethernet1/3
i L2      10.1.12.0/24 [115/20] via 10.1.13.1, 00:00:30, Ethernet1/3
```

On PE1

```
PE1#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      2.2.2.2 [115/20] via 10.1.12.2, 00:01:39, Ethernet1/0
i L2      3.3.3.3 [115/20] via 10.1.13.3, 00:00:52, Ethernet1/3
```

On PE2

```
PE2#Sh ip route isis | i L2

       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
i L2      1.1.1.1 [115/20] via 10.1.12.1, 00:01:09, Ethernet1/0
i L2      3.3.3.3 [115/30] via 10.1.12.1, 00:00:19, Ethernet1/0
i L2      10.1.13.0/24 [115/20] via 10.1.12.1, 00:01:09, Ethernet1/0

PE2#Ping 3.3.3.3 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 3.3.3.3, timeout is 2 seconds:
Packet sent with a source address of 2.2.2.2
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/48/56 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Configure the backbone to support MPLS VPN for IPv6 connectivity between PE2 and PE3.

On PE2

```
PE2(config)#router bgp 123
PE2(config-router)#No bgp default ipv4

PE2(config-router)#neighbor 3.3.3.3 remot 123
PE2(config-router)#neighbor 3.3.3.3 update lo0

PE2(config-router)#address vpnv6 unicast
PE2(config-router-af)#neighbor 3.3.3.3 activate
```

The VPNv6 address-family is used to for MPLS VPN Backbone connectivity when supporting Route Distinguisher's and Route Targets. Since labels are not supported in LDP for IPv6 routes currently, IPv4 is used to establish connectivity for path setup.

On PE3

```
PE3(config)#router bgp 123
PE3(config-router)#No bgp default ipv4
PE3(config-router)#neighbor 2.2.2.2 remot 123
PE3(config-router)#neighbor 2.2.2.2 update lo0

PE3(config-router)#address vpnv6 unicast
PE3(config-router-af)#neighbor 2.2.2.2 activate
```

The "send-community" and "next-hop-self-self" options do need to be configured as they are defaults but it is recommended to configure them for best practice.

```
PE3#Sh ip bgp vpnv6 uni all summary | b Neigh

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.2         4          123       4       4        1    0    0 00:00:52        0

PE3#Sh ip bgp vpnv6 uni all neighbors | s Neighbor capabilities

  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Multisession Capability: advertised and received
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On PE1

```
PE1(config)#mpls label range 100 199
PE1(config)#mpls label prot ldp
PE1(config)#mpls ldp router-id lo0

PE1(config)#int e1/0
PE1(config-if)#mpls ip
PE1(config-if)#int e1/3
PE1(config-if)#mpls ip
```

On PE2

```
PE2(config)#mpls label range 200 299
PE2(config)#mpls label prot ldp
PE2(config)#mpls ldp router-id lo0

PE2(config)#int e1/0
PE2(config-if)#mpls ip
```

On PE3

```
PE3(config)#mpls label range 300 399
PE3(config)#mpls label prot ldp
PE3(config)#mpls ldp router-id lo0

PE3(config)#int e1/3
PE3(config-if)#mpls ip
```

Task 4

Configure VRF's on the PE Routers using a name of "TST". They should ONLY support the IPv6 protocol. The RD for both VRF's should be 1:20 on PE2 and 1:30 on PE3. Both VRF's should use 1:1000 to identify the VPN.

On PE2

```
PE2(config)vrf definition TST
PE2(config-vrf)#RD 1:10
PE2(config-vrf)#address-family ipv6
PE2(config-vrf-af)#route-target both 1:1000
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
PE2(config-vrf-af)#int g0/0
PE2(config-if)#vrf forwarding TST

% Interface GigabitEthernet0/0 IPv4 disabled and address(es) removed due
to disabling VRF TST
% Interface GigabitEthernet0/0 IPv6 disabled and address(es) removed due
to disabling VRF TST
```

The First "%" indicates the loss of the link local address. The second "%" indicates the loss of the global IPv6 address. Both addresses are removed as a result of applying the VRF to the interface and must be configured again.

```
PE2(config)#int g0/0
PE2(config-if)#ipv6 address FE80::2 link-local
PE2(config-if)#ipv6 address 10:10:20::2/64
```

On PE3

```
PE3(config)vrf definition TST
PE3(config-vrf)#RD 1:20
PE3(config-vrf)#address-family ipv6
PE3(config-vrf-af)#route-target both 1:1000

PE3(config-vrf-af)#int g0/0
PE3(config-if)#vrf forwarding TST
PE3(config-if)#ipv6 address FE80::3 link-local
PE3(config-if)#ipv6 address 10:10:30::3/64
```

The "vrf definition" command is used istead of the standard "ip vrf" command due to its ability to support the IPv6 address-family.

Task 5

The CE routers should be configured to allow connectivity between the loopback interfaces. A BGP AS number between 0 – 65536 is not allowed on the CE's.

On CE2

```
CE2(config)#router bgp 65537

*Jun 29 14:00:35.887: %BGP-4-NORTRID: BGP could not pick a router-id.
Please configure manually.
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The 4 byte ASN style is supported in the IOS relese used in this workbook (12.2.33SRE). This allows a maximum BGP number of 4,294,967,295. A router-id must be chosen when using the 4 byte ASN format. It will still be a 32 bit IPv4 prefix that does NOT have to exist on the router.

```
CE2(config-router)#bgp router-id 22.22.22.22
CE2(config-router)#no bgp default ipv4
CE2(config-router)#neighbor 10:10:20::2 remot 123

CE2(config-router)#address-family ipv6 unicast
CE2(config-router-af)#neighbor 10:10:20::2 activate
CE2(config-router-af)#netw 22:22:22::22/128
```

On PE2

```
PE2(config)#router bgp 123
PE2(config-router)#address-family ipv6 vrf TST
PE2(config-router-af)#neighbor 10:10:20::10 remote-as 65537
PE2(config-router-af)#neighbor 10:10:20::10 as-override
```

The AS-Override command is required if we use the same AS on both sides of the connection.

```
PE2#Sh ip bgp vpnv6 uni vrf TST summary | b Neighbor

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10:10:20::10    4        65537       7       6        2    0    0 00:01:51        1

PE2#Sh ip bgp vpnv6 uni vrf TST | b Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf TST)
*> 22:22:22::22/128 10:10:20::10             0             0 65537 i
```

On CE3

```
CE3(config)#router bgp 1.1
CE3(config-router)#bgp router-id 33.33.33.33
CE3(config-router)#No bgp default ipv4
CE3(config-router)#neighbor 10:10:30::3 remot 123

CE3(config-router)#address ipv6 uni
CE3(config-router-af)#neighbor 10:10:30::3 activate
CE3(config-router-af)#network 33:33:33::33/128

CE3#Sh run | i router bgp
router bgp 65537
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The 4 byte ASN can also be configured in the (#.#) format. The 1.1 would be:
00000000 00000001. 00000000 00000001
65536 + 1 = 65537
The running configuration also reveals the decimal format of the value.

On PE3

```
PE3(config)#router bgp 123
PE3(config-router)#address-family ipv6 vrf TST
PE3(config-router-af)#neighbor 10:10:30::10 remote-as 1.1
PE3(config-router-af)#neighbor 10:10:20::10 as-override


PE3#Sh ip bgp vpnv6 uni vrf TST summ | b Neigh

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10:10:30::10    4        65537       6       7        4    0    0 00:02:07        1

PE3#Sh ip bgp vpnv6 uni vrf TST | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:20 (default for vrf TST)
*>i22:22:22::22/128 ::FFFF:2.2.2.2           0    100      0 65537 i
*> 33:33:33::33/128 10:10:30::10             0             0 65537 i
```

On CE3

```
CE3#Ping 22:22:22::22 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 22:22:22::22, timeout is 2 seconds:
Packet sent with a source address of 33:33:33::33
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/16/40 ms
```

Verifying the Label Creation (Control Plane):

On CE2

```
CE2#Sh ip bgp ipv6 uni | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
*> 22:22:22::22/128 ::                       0         32768 i
*> 33:33:33::33/128 10:10:20::2                            0 123 123 i
```

CE2 owns the local 22:22:22::22 prefix and sends it in a BGP update to PE2.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On PE2

```
PE2#Sh ip bgp vpnv6 unicast vrf TST labels

   Network          Next Hop      In label/Out label
Route Distinguisher: 1:10 (TST)
   22:22:22::22/128 10:10:20::10    204/nolabel
   33:33:33::33/128 ::FFFF:3.3.3.3  nolabel/303
```

PE2 receives the update from CE 2 nd sets the next hop toward the originator. An inner label (204) is created. If a prefix were to arrive with label "204" it would be removed and sent to CE2. The MP-BGP label is sent to PE3 in a BGP Update.

On PE3

```
PE3#Sh ip bgp vpnv6 unicast vrf TST labels

   Network          Next Hop      In label/Out label
Route Distinguisher: 1:20 (TST)
   22:22:22::22/128 ::FFFF:2.2.2.2   nolabel/204
   33:33:33::33/128 10:10:30::10     303/nolabel
```

PE3 receives the BGP update from PE2 and maps the label. If it were to receive a packet that was unlabeled toward the 22:22:22::22/128 destination, it would label the packet initially with "204" as the inner tag. The prefix is propagated tp CE3 in a BGP Update.

On CE3

```
CE3#Sh ip bgp ipv6 uni | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
*> 22:22:22::22/128 10:10:30::3                              0 123 123 i
*> 33:33:33::33/128 ::                       0         32768 i
```

CE3 receives the update and sets the next hop to be PE3. Notice how the AS Path has been modified to allow it to receive the update in the first place.

Verifying the Label Creation (Data Plane):

On PE3

```
PE3#Sh mpls forwarding vrf TST 22:22:22::22/128 detail

Local      Outgoing   Prefix           Bytes Label   Outgoing   Next Hop
Label      Label      or Tunnel Id     Switched      interface
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
None       204        22:22:22::22/128[V]    \
                                  0             Et1/3       10.1.13.1
      MAC/Encaps=14/22, MRU=1496, Label Stack{100 204}
      CA00067A001FCA02067A001F8847 00064000000CC000
      VPN route: TST
      No output feature configured
```

When PE3 receives a packet, it will impose the inner tag (204) with a Bottom of the Stack set to a value of "1". The outermost tag of "100" will also be imposed but the Bottom of the Stack bit will not be set.

```
PE3#Sh mpls forwarding-table 2.2.2.2

Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
302        100        2.2.2.2/32       0             Et1/3       10.1.13.1
```

Tag "100" refers to the path destination. In this case, it is the IPv4 destination loopback on PE2. (2.2.2.2/32)

On PE1

```
PE1#Sh mpls forwarding-table 2.2.2.2

Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
100        Pop Label  2.2.2.2/32       26520         Et1/0       10.1.12.2
```

PE1 will receive label "100" and will perform normal LSR function by label swapping 100 to label "3" (Pop Label).

On PE2

```
PE2#Sh mpls forwarding-table vrf TST

Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
204        No Label   22:22:22::22/128[V]    \
                                  570           Gi0/0       FE80::22
```

PE2 receives the packet with no outer label and only the inner label is exposed with a value of "204". PE2 removes this tag and forwards it out of the interface asscociated with the VPN "TST".

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

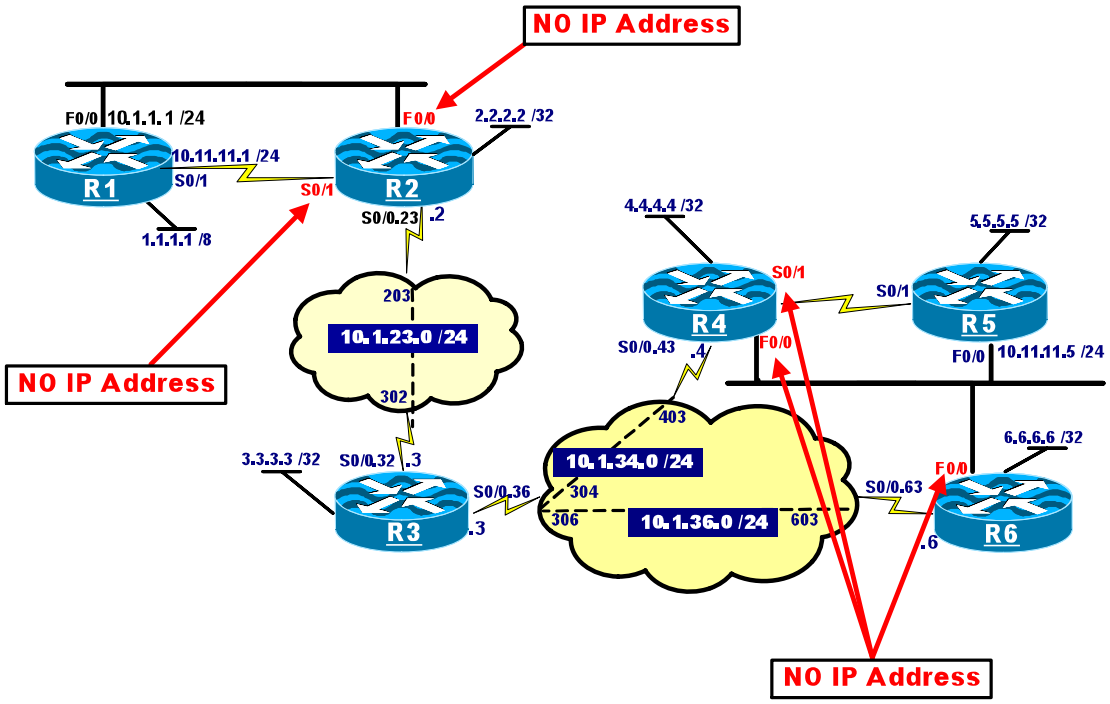
Paul Negron
CCIE #14856
SP

Layer 2 VPNs

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

LAB 1
Configuring MPLS L2 VPN's



IP Addressing:

.

Router	Interface	IP address
R1	Lo0 F0/0 S0/1	1.1.1.1/ 8 10.1.1.1 10.11.11.1
R2	Lo0 F0/0 S0/0.23 S0/1	2.2.2.2 /32 no ip address 10.1.23.2 /24 no ip address
R3	Lo0 S0/0.32 S0/0.34 S0/0.36	3.3.3.3/ 32 10.1.23.3 /24 10.1.34.3 /24 10.1.36.3/ 24
R4	Lo0	4.4.4.4 /32

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

	F0/0 S0/0.43 S0/1	no ip address 10.1.34.4/ 24 no ip address
R5	Lo0 F0/0 S0/1	5.5.5.5 /32 10.1.1.5/ 24 10.11.11.5/ 24
R6	Lo0 F0/0 S0/0.63	6.6.6.6/32 no ip address 10.1.36.6/ 24

Lab Setup:

- Routers R4, R5 and R6 are in VLAN 456.

Task 1

Configure OSPF "1"on the core devices R2, R3, R4 and R6.

On R2

```
R2(config)#router ospf 1
R2(config-if)#netw 0.0.0.0 0.0.0.0 area 0
```

On R3

```
R3(config)#router ospf 1
R3(config-if)#netw 0.0.0.0 0.0.0.0 area 0
```

On R4

```
R2(config)#router ospf 1
R2(config-if)#netw 0.0.0.0 0.0.0.0 area 0
```

On R6

```
R2(config)#router ospf 1
R2(config-if)#netw 0.0.0.0 0.0.0.0 area 0
```

Verify the OSPF Configuration:

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Sh ip route ospf | i O

O       3.3.3.3 [110/65] via 10.1.23.3, 00:00:10, Serial0/0.23
O       4.4.4.4 [110/129] via 10.1.23.3, 00:00:10, Serial0/0.23
O       6.6.6.6 [110/129] via 10.1.23.3, 00:00:10, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:00:10, Serial0/0.23
O       10.1.36.0 [110/128] via 10.1.23.3, 00:00:10, Serial0/0.23
```

On R3

```
R3#Sh ip route ospf | i O

O       2.2.2.2 [110/65] via 10.1.23.2, 00:00:34, Serial0/0.32
O       4.4.4.4 [110/65] via 10.1.34.4, 00:00:34, Serial0/0.34
O       6.6.6.6 [110/65] via 10.1.36.6, 00:00:34, Serial0/0.36
```

On R4

```
R4#Sh ip route ospf | i O

O       2.2.2.2 [110/129] via 10.1.34.3, 00:00:48, Serial0/0.43
O       3.3.3.3 [110/65] via 10.1.34.3, 00:00:58, Serial0/0.43
O       6.6.6.6 [110/129] via 10.1.34.3, 00:00:48, Serial0/0.43
O       10.1.23.0 [110/128] via 10.1.34.3, 00:00:58, Serial0/0.43
O       10.1.36.0 [110/128] via 10.1.34.3, 00:00:58, Serial0/0.43
```

On R6

```
R6#Sh ip route ospf | i O

O       2.2.2.2 [110/129] via 10.1.36.3, 00:01:21, Serial0/0.63
O       3.3.3.3 [110/65] via 10.1.36.3, 00:01:21, Serial0/0.63
O       4.4.4.4 [110/129] via 10.1.36.3, 00:01:21, Serial0/0.63
O       10.1.23.0 [110/128] via 10.1.36.3, 00:01:21, Serial0/0.63
O       10.1.34.0 [110/128] via 10.1.36.3, 00:01:21, Serial0/0.63
```

Task 2

Configure Layer 2 connectivity between R1 and R5 without the use of Bridge Groups. The path through R2>R3>R4>R5 must be used to accomplish the task.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config)#int f0/0
R2(config-if)#xconnect 4.4.4.4 204 encapsulation mpls
R2(config-if-xconn)#end
```

On R4

```
R4(config)#int f0/0
R4(config-if)#xconnect 2.2.2.2 204 encapsulation mpls
R4(config-if-xconn)#end
```

The xconnect command used on the F0/0 interface on both routers is used to create a bridged connection with the destination specified. The command is broken down with an xconnect keyword followed by the peering router address and the unique virtual circuit ID (VCID). The VCID must match on both ends of the connection. The encapsulation can be L2TPv2, L2TPv3 or MPLS. There must be a unique address per router for each xconnect.

On R2

```
R2#Show xconnect all

Legend: XC ST=Xconnect State, S1=Segment1 State, S2=Segment2 State
UP=Up, DN=Down, AD=Admin Down, IA=Inactive, NH=No Hardware
XC ST  Segment 1                  S1 Segment 2                  S2
------+-----------------------------+--+-----------------------------+--
DN     ac   Fa0/0(Ethernet)         AD mpls 4.4.4.4:204           DN
```

An MPLS Layer 2 Pseudowire has 2 segments. Segment 1 (S1) is for the Customer facing port. Segment 2(S2) is relating to the core configuration. The output of the show command reveals that Segment 1 of the connection is Administratively Shutdown.

On R1,R2,R4 and R5

```
RX(config)#int f0/0
RX(config-if)#NO shut
```

On R2

```
R2#Sh xconnect peer 4.4.4.4 vcid 204 detail

Legend: XC ST=Xconnect State, S1=Segment1 State, S2=Segment2 State
UP=Up, DN=Down, AD=Admin Down, IA=Inactive, NH=No Hardware
XC ST  Segment 1                  S1 Segment 2                  S2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
------+-----------------------------+--+-----------------------------+--
DN    ac   Fa0/0(Ethernet)          UP mpls 4.4.4.4:204             DN

          Interworking: none              Local VC label 21
                                          Remote VC label 21
                                          pw-class:
```

The physical ports have been activated but the XConnect still remains down due to Segment 2(S2) not being configured yet.

Setup the MPLS Core with LDP:

On R2

```
R2(config)#int s0/0.23
R2(config-subif)#mpls ip
```

On R3

```
R3(config)#int s0/0.32
R3(config-subif)#mpls ip
R3(config-subif)#int s0/0.34
R3(config-subif)#mpls ip
```

On R4

```
R4(config)#int s0/0.43
R4(config-subif)#mpls ip
```

LDP was chosen to provide Core MPLS connectivity although Traffic Engineering could have also been used. This lays the Segment 2 foundation needed for the pseudowire to operate.

Verify the Xconnect pseudowire:

On R4

```
R4#Show xconnect peer 2.2.2.2 vcid 204

Legend: XC ST=Xconnect State, S1=Segment1 State, S2=Segment2 State
UP=Up, DN=Down, AD=Admin Down, IA=Inactive, NH=No Hardware
XC    ST  Segment 1                 S1 Segment 2                  S2
------+-----------------------------+--+-----------------------------+--
UP    ac  Fa0/0(Ethernet)           UP mpls 2.2.2.2:204             UP
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The S2 section has now been configured and verified to be working properly with state (ST) of active (ac).

On R2

```
R2#Sh mpls ldp neighbo 4.4.4.4 | b LDP discovery sources:

        LDP discovery sources:
          Targeted Hello 2.2.2.2 -> 4.4.4.4, active, passive
      Addresses bound to peer LDP Ident:
        10.1.34.4       4.4.4.4
```

The peer exchange is unicast peering. Targeted hellos are being sent toward the other end of the peering session, which in this case is R4. The "active, passive" state of the targeted hello states that this connection has initiated hellos from the local side and recognizes that the far end has initiated hellos as well.

```
R2#Sh mpls forwarding-table

Local  Outgoing    Prefix           Bytes Label   Outgoing    Next Hop
Label  Label or VC or Tunnel Id     Switched      interface
16     Pop Label   3.3.3.3/32       0             Se1/0.23    point2point
17     17          4.4.4.4/32       0             Se1/0.23    point2point
18     18          6.6.6.6/32       0             Se1/0.23    point2point
19     Pop Label   10.1.34.0/24     0             Se1/0.23    point2point
20     Pop Label   10.1.36.0/24     0             Se1/0.23    point2point
21                 l2ckt(204)       0             none        point2point
```

The local label assigned to the router is "21" in this case. The label given to your router may be different. Notice that the label is listed as having a local tag with no remote label. The prefix is listed as being a Layer 2 circuit (l2ckt) with a vcid of 204 and will be mapped to Segment 2 which has it's own label number and process.

```
R2#Sh xconnect peer 4.4.4.4 vcid 204 detail

Legend: XC ST=Xconnect State, S1=Segment1 State, S2=Segment2 State
UP=Up, DN=Down, AD=Admin Down, IA=Inactive, NH=No Hardware
XC ST  Segment 1                  S1 Segment 2                  S2
------+-----------------------------+--+-----------------------------+--
UP     ac   Fa0/0(Ethernet)         UP mpls 4.4.4.4:204           UP
          Interworking: none            Local VC label 21
                                        Remote VC label 21
                                        pw-class:
```

The local label can also be viewed in the "show xconnect" command as well as the label assigned to the

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

remote peer. The label mapping is always from local to remote.

On R1

```
R1#Sh cdp n

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
R2               Fas 0/0           122        R           3725      Fas 0/0
R5               Fas 0/0           126        R S I       3725      Fas 0/0
```

CDP as well as all layer 2 protocols are being tunneled through in addition to the normal CDP packet from the directly connected router.

Configure Layer 3 information on R1 and R5:

On R1

```
R1(config)#router eigrp 1
R1(config-router)#NO auto
R1(config-router)#netw 1.0.0.0
R1(config-router)#netw 10.0.0.0
```

On R5

```
R5(config-if)#router eigrp 1
R5(config-router)#NO auto
R5(config-router)#netw 5.0.0.0
R5(config-router)#netw 10.0.0.0
```

Verify Connectivity:

On R1

```
R1#Sh ip route eigrp

D    5.0.0.0/8 [90/409600] via 10.1.1.5, 00:00:04, FastEthernet0/0

R1#Ping 5.5.5.5 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Packet sent with a source of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/30/64 ms
```

The routes can be viewed and reached transparently through the pseudowire connection.

Task 3

Configure a backup connection for the primary pseudowire connection created in Task 1. This connection should take the following path: R2>R3>R6>R5. The backup peer should not be used automatically in the event of a failure.

On R2

```
R2(config)#int f0/0
R2(config-if)#xconnect 4.4.4.4 204 encapsulation mpls
R2(config-if-xconn)#backup peer 6.6.6.6 206
R2(config-if-xconn)#backup delay 0 never
R2(config-if-xconn)#end
```

The backup connection is created as an option under the primary tunnel created in Task 1. The "backup delay" is how long it will wait to trigger to the backup tunnel, "0" with an option of never will never switch over until manually configured.

On R3

```
R3(config)#int s0/0.36 p
R3(config-subif)#mpls ip
```

On R6

```
R6(config)#int s0/0.63
R6(config-subif)#mpls ip

R6(config)#int f0/0
R6(config)#NO shut
R6(config-if)#xconnect 2.2.2.2 206 encapsulation mpls
R6(config-if-xconn)#end
```

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Sh xconnect all

Legend: XC ST=Xconnect State, S1=Segment1 State, S2=Segment2 State
UP=Up, DN=Down, AD=Admin Down, IA=Inactive, NH=No Hardware

XC ST  Segment 1                    S1 Segment 2                    S2
------+-------------------------------+--+-------------------------------+--
UP pri ac   Fa0/0(Ethernet)           UP mpls 4.4.4.4:204             UP
IA sec ac   Fa0/0(Ethernet)           UP mpls 6.6.6.6:206             DN
```

Notice how the primary tabe is "UP" on both Segments while the backup remains down on S2 until it is time to be used.

On R2

```
R2#xconnect backup force-switchover peer 6.6.6.6 206
```

We trigger the manual switchover.

```
R2#Sh xconnect all

Legend: XC ST=Xconnect State, S1=Segment1 State, S2=Segment2 State
UP=Up, DN=Down, AD=Admin Down, IA=Inactive, NH=No Hardware
XC ST  Segment 1                    S1 Segment 2                    S2
------+-------------------------------+--+-------------------------------+--
IA pri ac   Fa0/0(Ethernet)           UP mpls 4.4.4.4:204             DN
UP sec ac   Fa0/0(Ethernet)           UP mpls 6.6.6.6:206             UP
```

The backup tunnel is now "UP" and is operational due to the manual switchover.

On R1

```
R1#Ping 5.5.5.5 source 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 16/30/64 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 4

Configure a third xconnect between R1 and R5 using the respective Serial interface. HDLC should be used as the encapsulation on the physical link and provide a second EIGRP peer successfully.

On R1

```
R1(config)#int s0/0
R1(config-if)#No shut
```

On R2

```
R2(config)#int s0/1
R2(config-if)#No shut
R2(config-if)#xconnect 4.4.4.4 222 encapsulation mpls
R2(config-if-xconn)#end
```

On R4

```
R4(config)#int s0/1
R4(config-if)#No shut
R4(config-if)#xconnect 2.2.2.2 222 encapsulation mpls
R4(config-if-xconn)#end
```

The entire Serial interface will tunnel ALL layer 2 across the cloud.

On R5

```
R5(config)#int s0/1
R5(config-if)#No shut

R5#Sh ip ei n

IP-EIGRP neighbors for process 1
H   Address                 Interface       Hold Uptime   SRTT      RTO  Q  Seq
                                            (sec)         (ms)      Cnt Num
1   10.11.11.1              Se0/1           12 00:00:34   30        200  0  13
0   10.1.1.1                Fa0/0           14 00:20:19   43        258  0  11
```

The secondary peer should come up due to the EIGRP configuration performed in Task 2.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

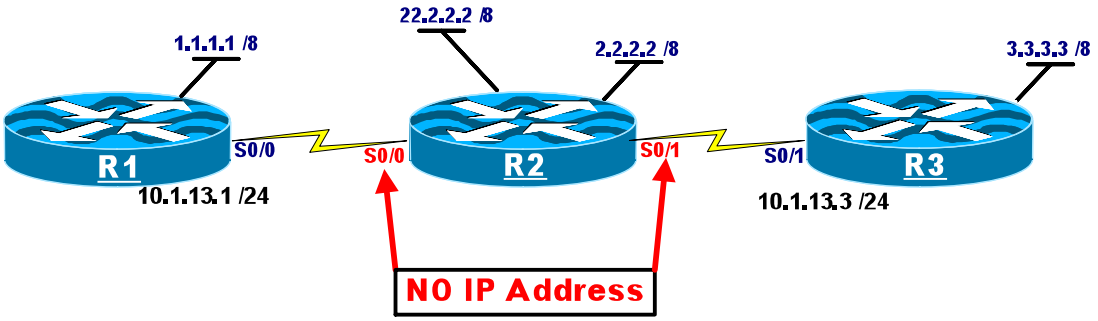
Task 5

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 2 –L2TPv3 VPN



IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0	1.1.1.1/ 8 10.1.13.1/ 24
R2	Lo0 Lo1 S0/0 S0/1	2.2.2.2 /8 22.2.2.2/ 8 no ip address no ip address
R3	Lo0 S0/1	3.3.3.3/ 8 10.1.13.3/ 24

Task 1

Configure a Layer 2 VPN between R1 and R3.

On R1

```
R1(config)#interface s0/0
R1(config-if)#NO shut
R1(config-if)#end
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#interface s0/1
R1(config-if)#NO shut

R1(config-if)#end
```

On R2

```
R2(config)#pseudowire-class TSTR1
R2(config-pw-class)#encapsulation L2tpv3
R2(config-pw-class)#protocol L2tpv3
R2(config-pw-class)#ip local int lo0


R2(config)#pseudowire-class TSTR3
R2(config-pw-class)#encapsulation L2tpv3
R2(config-pw-class)#protocol L2tpv3
R2(config-pw-class)#ip local int lo1
```

Unlike using the mpls encapsulation, the L2tpv3 service requires the use of the source interface. Therefore, this back-to-back configuration is supported.

```
R2(config)#int s0/0
R2(config-if)#NO shut
R2(config-if)#xconnect 22.2.2.2 222 pw-class TSTR1
R2(config-if-xconn)#exit
```

The pseudowire class is applied to the interface instead of specifying the encapsultion directly.

```
R2(config-if)#int s0/1
R2(config-if)#NO shut
R2(config-if)#xconnect 2.2.2.2 222 pw-class TSTR3
R2(config-if-xconn)#end
```

On R1

```
R1#Sh cdp n

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater

Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
R2               Ser 0/0           126           R S I    3725      Ser 0/0
R3               Ser 0/0           161           R S I    3725      Ser 0/1

R1#Ping 10.1.13.3
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.3, timeout is 2 seconds:

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/24/64 ms
```

Task 2

The Layer 2 VPN peers should be used with a username of R1 and a password of cisco on the R1 side of the connection and a Username of R3 on the R3 side of the connection. Fragments should not be permitted across the VPN. The control messages should be encrypted with the strongest encryption method and password of CCIE.

On R2

```
R2(config)#l2tp-class TESTR1
R2(config-l2tp-class)#hostname R1
R2(config-l2tp-class)#authentication
R2(config-l2tp-class)#password cisco
R2(config-l2tp-class)#digest secret CCIE hash SHA1
R2(config-l2tp-class)#exit

R2(config)#l2tp-class TESTR3
R2(config-l2tp-class)#hostname R3
R2(config-l2tp-class)#authentication
R2(config-l2tp-class)#password cisco
R2(config-l2tp-class)#digest secret CCIE hash SHA1
R2(config-l2tp-class)#exit
```

A l2tp-class template is used to modify options dealing with the Layer 2 connection messages and some of the peering characteristics. The followig options are available in the l2tp-class:

***Hostname*: the hostname used by this router in the connection. If not specified, the hostname of the router is used.**

***Authentication/ Password:* specifies authentication is to be used and the password that is used to authenticate the peering session.**

***Digest secret:* the encryption used for he control messages between peers. The hash can be configurd to use MD5 and SHA1.**

```
R2(config)#pseudowire-class TSTR1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-pw-class)#protocol L2tpv3 TESTR1
```

L2TP class changes are not allowed on L2TP pw-class with Xconnects

The L2tp-class cannot be modified while the xconnect exist.

Remove the Xconnects:

```
R2(config)#int s0/0
R2(config-if)#NO xconnect 22.2.2.2 222
R2(config-if)#int s0/1
R2(config-if)#NO xconnect 2.2.2.2 222

R2(config)#pseudowire-class TSTR1
R2(config-pw-class)#protocol L2tpv3 TESTR1

R2(config)#pseudowire-class TSTR3
R2(config-pw-class)#protocol L2tpv3 TESTR3
```

The pseudowire-class has information that deals with the encapsuation between peer as well as options dealing with the Don't Fragment bit (DF), ip tos, and ip ttl of the tunnel. The l2tp-class is applied to the pseudowire class via the "protocol" option. The local interface must be specified. Two different loopbacks are needed, as the same interface cannot be used for two different pseudowire tunnels.

```
R2(config)#int s0/0
R2(config-if)#xconnect 22.2.2.2 222 pw-class TSTR1
R2(config-if-xconn)#exit

R2(config-if)#int s0/1
R2(config-if)#xconnect 2.2.2.2 222 pw-class TSTR3
R2(config-if-xconn)#end
```

On R1

```
R1#Ping 10.1.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 8/24/64 ms
```

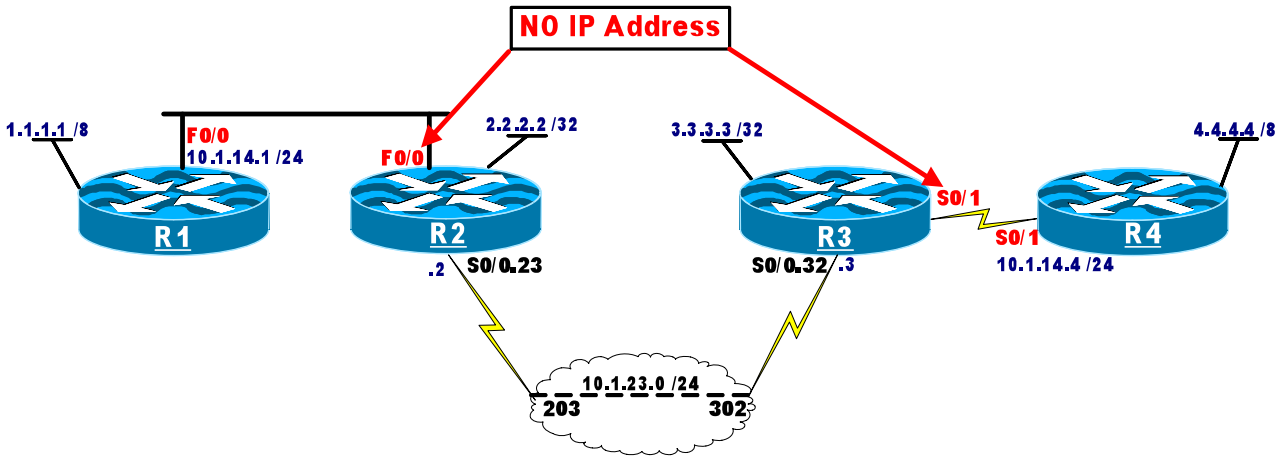

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 3 –L2TPv3-Interworking



IP Addressing:

Router	Interface	IP address
R1	Lo0 F0/0	1.1.1.1/ 8 10.1.14.1/ 24
R2	Lo0 F0/0 S0/0.23	2.2.2.2 /32 no ip address 10.1.23.2/ 24
R3	Lo0 S0/0.32 S0/1	3.3.3.3/ 32 10.1.23.3/ 24 no ip address
R4	Lo0	4.4.4.4/8

Task 1

Configure OSPF with a process of "1" as the IGP between R2 and R3.

On R2

```
R2(config)#router ospf 1
R2(config-router)#netw 0.0.0.0 255.255.255.255 area 0
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#router ospf 1
R3(config-router)#netw 0.0.0.0 255.255.255.255 area 0
```

Verifying the Configuration:

On R2

```
R2#Sh ip route ospf | i O

O        3.3.3.3 [110/65] via 10.1.23.3, 00:06:51, Serial0/0.23
```

On R3

```
R3#Sh ip route ospf | i O

O        2.2.2.2 [110/65] via 10.1.23.2, 00:05:23, Serial0/0.32
```

Task 2

Configure a Layer 2 VPN between R1 and R4 using 2 different framing types. Do not allow the session to fragment packets.

On R1

```
R1(config)#interface f0/0
R1(config-if)#NO shut
```

On R2

```
R2(config)#pseudowire-class TST
R2(config-pw-class)#encapsulation L2tpv3
R2(config-pw-class)#ip local int lo0
R2(config-pw-class)#ip dfbit set
R2(config-pw-class)#interworking ip
```

The "df" bit is set, which will not allow fragmentation. The lack of fragmentation does not allow us to use sequencing of packets across the link.

Interworking is the keyword necessary to have framing types terminate locally on this router. In this case, the framing type is Ethernet on this side and it is Frame-Relay on R3.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-pw-class)#int f0/0
R2(config)#NO shut

R2(config-if)#xconnect 3.3.3.3 203 pw-class TST
R2(config-if-xconn)#exit

R2(config)#int s0/0.23 p
R2(config-subif)#mpls ip
```

On R3

```
R3(config)#frame-relay switching
```

Since the router is technically performing dlci switching and also is responsible for dlci allocation for the customer, we must enable frame-relay switching to allow it to behave like one.

```
R3(config)#pseudowire-class TST
R3(config-pw-class)#encapsulation L2tpv3
R3(config-pw-class)#ip local int lo0
R3(config-pw-class)#ip dfbit set
R3(config-pw-class)#interworking ip

R3(config)#int s0/1
R3(config-if)#encap frame-relay
R3(config-if)#frame-relay intf-type dce
R3(config-if)#NO shut
R3(config-if)#exit
```

The frame relay switch also provides clocking for the layer 2 DCE as well, just like a frame switch.

```
R3(config)#connect TSTR4 Serial 0/1 401 l2transport
R3(config-fr-pw-switching)#xconnect 2.2.2.2 203 pw-class TST
R3(config-xconn)#exit
```

The name of the connection "TSTR4" is locally significant.

```
R3(config-fr-pw-switching)#int s0/0.32
R3(config-subif)#mpls ip
```

On R4

```
R4(config)#int s0/1
R4(config-if)#encap frame-relay
R4(config-if)#NO shut
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-if)#int s0/1.41 p
R4(config-subif)#ip add 10.1.14.4 255.255.255.0
R4(config-subif)#frame inter 401
```

Verifying the Configuration:

On R2

```
R2#Show l2tun tunnel transport

L2TP Tunnel Information Total tunnels 1 sessions 1

LocTunID  Type Prot   Local Address  Port   Remote Address  Port
32551     IP   115    2.2.2.2        0      3.3.3.3         0
```

The L2 Tunnel has been formed using port 115 and uses a Local Identifer of 32551, which can be seen as a Remote Identifier on R3. The ID's may not match yours but both should be consistent.

On R3

```
R3#Sh xconnect peer 2.2.2.2 vcid 203 detail

Legend: XC ST=Xconnect State, S1=Segment1 State, S2=Segment2 State
UP=Up, DN=Down, AD=Admin Down, IA=Inactive, NH=No Hardware
XC ST  Segment 1                  S1 Segment 2                  S2
------+-------------------------------+--+-------------------------------+--
UP     ac  Se0/1:401(FR DLCI)         UP l2tp 2.2.2.2:203           UP
           Interworking: ip                Session ID: 7457
                                  Tunnel ID: 17066
                                  Protocol State: UP
                                  Remote Circuit State: UP
                                  pw-class: TST
```

Interworking has been noted as being IP. This translation takes place across two sides of a conection that do not match encapsulations. The IP packet is delivered into the correct Framing locally on the router.

```
R3#Sh l2tp

L2TP Tunnel and Session Information Total tunnels 1 sessions 1

LocTunID   RemTunID   Remote Name   State  Remote Address  Sessn L2TP Class/
                                                Count VPDN Group
17066      32551      R2            est    2.2.2.2         1     l2tp_default_cl
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
LocID      RemID      TunID      Username, Intf/      State  Last Chg Uniq ID
                                 Vcid, Circuit
7457       59753      17066      203, Se0/1:401       est    00:04:55 1
```

The local ID and remote ID remain consistent with the output seen earlier. The use of the defaut pseudowire class was used instead of using a custom configuration.

On R1

```
R1#Sh cdp n

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
Device ID        Local Intrfce     Holdtme    Capability  Platform  Port ID
```

CDP updates may be seen local with R2 but are not seen across the pseudowire due to the fact that this is not Layer 2 interworking. The layer 3 interworking should be fine.

```
R4#Ping 10.1.14.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 16/24/40 ms
```

And it is!!!

Task 3

The Packet should die on it's own after 50 hops. QOS should bump the session to "IMMEDIATE". Both ends should negotiate the MTU to be 1400 bytes.

On R2

```
R2#Sh l2tp session all | s ToS

  DF bit on, ToS reflect disabled, ToS value 0, TTL value 255

R2(config)#pseudowire-class TST
R2(config-pw-class)#ip pmtu max 1400
R2(config-pw-class)#ip tos value 32
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config-pw-class)#ip ttl 50
```

The path mtu has been set for 1400. The default value of "1500" is not shown in the output.

The tos value is set to equal an IP Precedence of '001' in the 3 high order bits of the octet.

The "ip ttl" command has been modified to 50 to have the packet die once the 51st HOP occurs.

On R3

```
R3(config)#pseudowire-class TST
R3(config-pw-class)#ip tos value 32
R3(config-pw-class)#ip pmtu max 1400
R3(config-pw-class)#ip ttl 50
```

Verifying the Parameters

On R2

```
R2#Sh l2tp session all | i 1400

  Session PMTU enabled, path MTU is 1400 bytes

R2#Sh l2tp session all | s ToS

  DF bit on, ToS reflect disabled, ToS value 32, TTL value 50
```

On R3

```
R3#Sh l2tp session all | i 1400

  Session PMTU enabled, path MTU is 1400 bytes

R3#Sh l2tp session all | s ToS

  DF bit on, ToS reflect disabled, ToS value 32, TTL value 50
```

All of the parameters have been changed and verified. A resetting was not needed for the connection to work immediately.

On R1

```
R1#Ping 10.1.14.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.14.4, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 36/48/56 ms
```

If the MTU parameter did not match, the ping would be unsuccessful.

Task 3

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

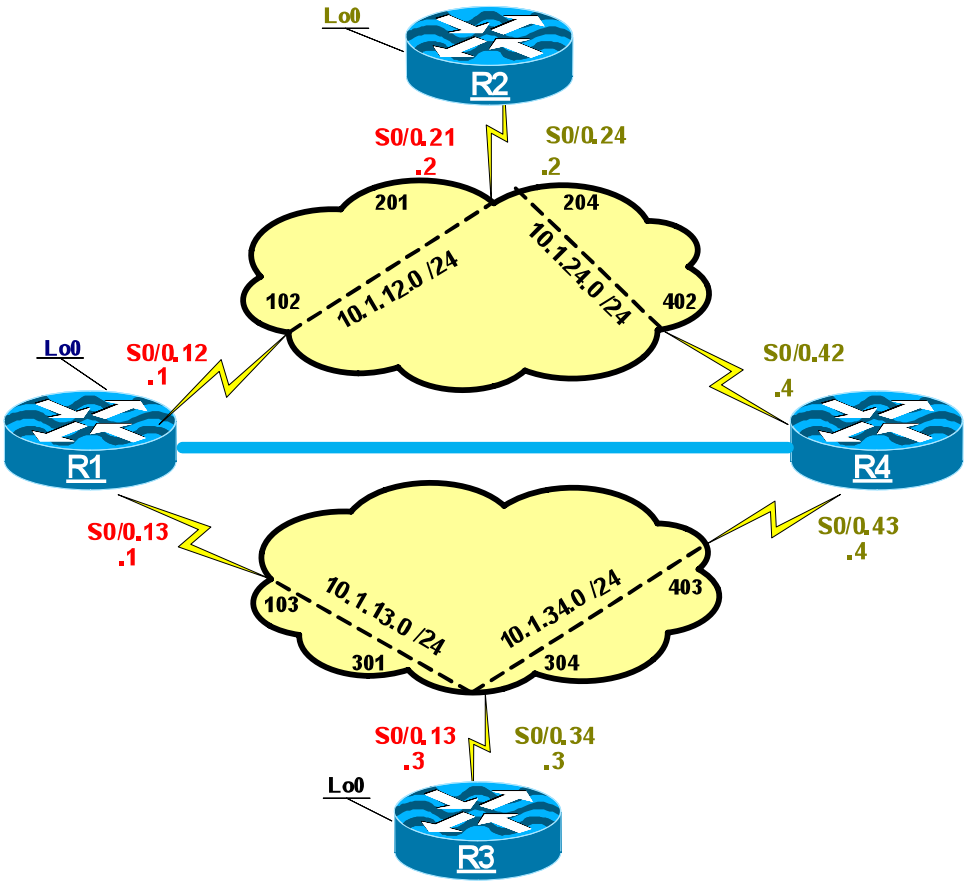
Paul Negron
CCIE #14856
SP

MPLS Traffic Engineering

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 1 – Configuring Static Tunnels



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12 S0/0.13 F0/0	1.1.1.1 /32 10.1.12.1 /24 10.1.13.1 /24 10.1.14.1/ 24
R2	Lo0 S0/0.21 S0/0.23	2.2.2.2 /32 10.1.12.2 /24 10.1.24.2 /24
R3	Lo0 S0/0.31 S0/0.34	3.3.3.3/ 32 10.1.13.3 /24 10.1.34.3 /24
R4	Lo0 F0/0 S0/0.42 S0/0.43	4.4.4.4 /32 10.1.14.4/ 24 10.1.24.4 /24 10.1.34.4 /24

Lab Setup:

The connections between R1 to R2, R1 to R3, R2 to R4, and R3 to R4 are to be setup in a Frame-Relay Point-to-Point manner

Task 1

Configure an OSPF routed core using process "1". Support mpls traffic engineering throughout the core topology.

Configure basic traffic Engineering on ALL Core Devices:

On R1

```
R1(config)#mpls traffic-eng tunnels
```
(This command allows MPLS-TE globally)

Configure on all core routers R1, R2, R3 and R4 with this command.

```
R1(config)#router ospf 1
R1(config)#network 0.0.0.0 0.0.0.0 area 0
R1(config-router)#mpls traffic-eng area 0
R1(config-router)#mpls traffic-eng router-id loopback 0
```

OSPF needs to recognize and use the OSPF Opaque LSA (type 10) to include the Traffic Engineering paths that will be configured. This command must also be configured on ALL core routers.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config)#interface s0/0.12
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth 60

R1(config)#interface f0/0
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth 60

R1(config-if)#interface s0/0.13
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth (we leave this option with the default)
```

ALL interfaces that may carry tunnel traffic at any time must be configured with the "mpls traffic-eng tunnels" command on the interface as well. If this command is not issued at this level, the router will not participate in MPLS Traffic Engineering path calculation (PCALC).

ALL interfaces that may carry tunnel traffic at any time must be configured with the "ip rsvp bandwidth" command on the interface that could be a potential path. The router will not participate in MPLS Traffic Engineering path calculation (PCALC) from the head-end of the tunnel if not configured. The default value, if not specified, will be 75% of the "bandwidth" command configured on the interface. 1.544 Mbps is the default bandwith of any of the serial interfaces. RSVP will use this value to subtract any requested bandwidth by a tunnel.

The three configurations: "mpls traffic-eng tunnels"(global), "mpls traffic-eng tunnels"(interface), and "ip rsvp bandwith" must all be used on all core interfaces to participate in MPLS Traffic Engineering, even if there will be no tunnels built from the router itself.

On R2

```
R2(config)#mpls traffic-eng tunnels

R2(config)#router ospf 1
R2(config)#network 0.0.0.0 0.0.0.0 area 0
R2(config-router)#mpls traffic-eng area 0
R2(config-router)#mpls traffic-eng router-id loopback 0

R2(config)#interface s0/0.21
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth

R2(config-if)#interface s0/0.24
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#mpls traffic-eng tunnels

R3(config)#router ospf 1
R3(config)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#mpls traffic-eng area 0
R3(config-router)#mpls traffic-eng router-id loopback 0

R3(config)# interface s0/0.31
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth

R3(config-if)#interface s0/0.34
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth
```

On R4

```
R4(config)#mpls traffic-eng tunnels

R4(config)#router ospf 1
R4(config)#network 0.0.0.0 0.0.0.0 area 0
R4(config-router)#mpls traffic-eng area 0
R4(config-router)#mpls traffic-eng router-id loopback 0

R4(config)#interface s0/0.42
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth

R4(config)#interface f0/0
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth

R4(config-if)#interface s0/0.43
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth
```

Verifying basic traffic Engineering:

On R1

```
R1#Sh mpls traffic-eng link-management igp-neighbors

Link ID::  Fa0/0
    Neighbor ID:  4.4.4.4 (area: ospf area 0, IP: 10.1.14.4)
Link ID::  Se0/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    Neighbor ID:  2.2.2.2 (area: ospf area 0, IP: 10.1.12.2)
Link ID::  Se0/0.13
    Neighbor ID:  3.3.3.3 (area: ospf area 0, IP: 10.1.13.3)
```

MPLS traffic engineering has been correctly configured under the core interfaces or we would not be able to see neighbors.

```
R1#Sh ip ospf database | be Type-10
```

Type-10 Opaque Link Area Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Opaque ID
1.0.0.0	1.1.1.1	183	0x80000002	0x004347	0
1.0.0.0	2.2.2.2	148	0x80000002	0x00CCB9	0
1.0.0.0	3.3.3.3	112	0x80000002	0x004932	0
1.0.0.0	4.4.4.4	97	0x80000002	0x00CCE4	0
1.0.0.1	1.1.1.1	31	0x80000001	0x0038B6	1
1.0.0.1	2.2.2.2	31	0x80000001	0x008945	1
1.0.0.1	3.3.3.3	108	0x80000002	0x001C3E	1
1.0.0.1	4.4.4.4	93	0x80000002	0x00A50F	1
1.0.0.2	1.1.1.1	23	0x80000001	0x00AE77	2
1.0.0.2	4.4.4.4	75	0x80000001	0x00ACD5	2

The type 10 opaque LSA is used to identify all of the LSA's learned from the MPLS traffic-engineering topology. The highlighted paths all share the same link ID's which represents the tunnel path and all of the LSA's learned from that path.

```
R1#Sh ip rsvp int
```

interface	allocated	i/f max	flow max	sub max
F0/0	0	60K	60K	0
Se0/0	0	0	0	0
Se0/0.12	0	60K	60K	0
Se0/0.13	0	1158K	1158K	0

RSVP has been configured under each of the core interfaces as well. The default-reserved bandwidth is 75% of the physical bandwidth. Default would be Se0/0.13 = 1544000 – 386000 (25%) = 1158K. Since we set the bandwidth to restrict to 60K on the other interfaces, it overrides the default.

Configure R1 s0/0.13 with "ip rsvp bandwidth 60" command before proceeding!!

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Configure MPLS-TE to establish basic connectivity from R1 to R4. R1 should not follow the default routing path. The path through R2 should be favored as a static path with the following attributes. Priority = 4, all traffic uses the tunnel to 4.4.4.4 and bandwidth = 30K.

Verify Current Routing Path:

On R1

```
R1#Trace 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4

  1 10.1.14.4 4 msec * 4 msec

R1#Sh ip cef exact-route 1.1.1.1 4.4.4.4

1.1.1.1         -> 4.4.4.4         : FastEthernet0/0 (next hop 10.1.14.4)
```

The current path is currently R1-R4 through the F0/0 interface.

Change the Current Routing Path:

```
R1(config)#interface tunnel 0
R1(config-if)#ip unnumbered loopback 0
```

A good practice to not waste an ip address on the tunnel interface.

```
R1(config-if)#tunnel destination 4.4.4.4
```

The destination of the tunnel = the traffic-eng router-id of the last hop router

```
R1(config-if)#tunnel mode mpls traffic-eng
```

The tunnel mode must be changed from the default point-to-point tunnel mode, to mpls traffic engineering to hav RSVP create labels.

```
R1(config-if)#tunnel mpls traffic-eng autoroute announce
```

This command ensures that all traffic is routed into the tunnel and announced into ospf updates.

```
R1(config-if)#tunnel mpls traffic-eng  priority 4 4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The tunnel priority will use the SETUP priority (the first number) and compare with another tunnel's HOLD priority (the second number in the command). Both numbers are usually set to be the same to avoid confusion but may be different to establish true priority.

```
R1(config-if)#tunnel mpls traffic-eng bandwidth 30
R1(config-if)#tunnel mpls traffic-eng path-option 1 explicit name SLOW_LINK
R1(config-if)#exit
```

The tunnel bandwidth of "30" is in kbps and is what the tunnel requires to be setup.
The tunnel should still remain down until the explicit path is created.

```
R1(config)#ip explicit-path name SLOW_LINK enable
R1(config)#next-address 10.1.12.2
R1(config)#next-address 10.1.24.4
```

The path should come up after the last hop is configured

Verify the Traffic Engineered path Information

On R1

```
R1#Sh mpls traffic-eng tunnels brief

Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 3600 seconds, next in 2149
seconds
    Periodic auto-bw collection:    disabled

TUNNEL NAME             DESTINATION     UP IF   DOWN IF  STATE/PROT
R1_t0                   4.4.4.4         -       Se0/0.12 up/up
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 0 (of 0) tails
```

Notice that this is the "head" of the tunnel. It does not represent any other tunnel as a midpoint or a tail. The LSP Tunnel and RSVP processes show a running state.

Forwarding should be enabled or the top two processes do not matter.

Periodic reoptimazation allows for the tunnels to recalculate the best paths automatically over a specific amount of time (3600 seconds). There will be another update in 2149 seconds.

Periodic auto-bw collection is currently disabled and is used when automatically attempting to reroute based on utilization. This will be covered in another section.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The tunnel name, "R1_t0", can sometimes be referenced in the command line. The destination of the tunnel is the TE Router-ID of the last hop TE router R4.

The "UP IF" is a term used to describe the Downstream peer (next hop router) from a Data Plane perspective. The "DOWN IF" is "Serial 0/0.12". The router output locks the field in this output to a certain size. In some cases, the interface name is too long for the space. The interface can be seen with the "show mpls traffic-eng tunnels" command.

The State and Protocol are treated as layer 1 and Layer 2 for the physical interface. Possible states are "*up/up*"- the physical line and the line protocol are operational.

***up/down*- the physical layer is good for the tunnel but the line protocol has been broken as in the example of a sub-interface or physical interface failure. The tunnel interface is fine from a logical and physical layer view.**

```
R1#Sh mpls traffic-eng tunnels

Name: R1_t0                           (Tunnel0) Destination: 4.4.4.4
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected

    path option 1, type explicit SLOW_LINK (Basis for Setup, path weight 128)

  Config Parameters:
    Bandwidth: 30       kbps (Global)  Priority: 4  4   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 30       bw-based
    auto-bw: disabled

  InLabel  :  -
  OutLabel : Serial0/0.12, 16
  RSVP Signalling Info:
       Src 1.1.1.1, Dst 4.4.4.4, Tun_Id 0, Tun_Instance 4
    RSVP Path Info:
      My Address: 1.1.1.1
      Explicit Route: 10.1.12.2 10.1.24.4 4.4.4.4
      Record Route:  NONE
      Tspec: ave rate=30 kbits, burst=1000 bytes, peak rate=30 kbits
    RSVP Resv Info:
      Record Route:  NONE
      Fspec: ave rate=30 kbits, burst=1000 bytes, peak rate=30 kbits
  History:
    Tunnel:

      Time since created: 2 hours, 17 minutes
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Time since path change: 2 hours, 15 minutes
Current LSP:
Uptime: 2 hours, 15 minutes

The name, destination have already been covered.

The "*Admin: up*", "*Oper: up*, *Path: Valid* and *Signalling*: *connected* fields, make up extremly useful ouput for troubleshooting purposes. By the end of the labs, you should be able to use the "Path valid" and "Signalling" output to its' fullest extent.

The "path option 1" describes the number of the path-option declared by configuration that the tunnel has chosen. "type Explicit SLOW_LINK" was the explicit map configured earlier. The "(basis for Setup, setup weight 128)" field is igp/te cost = 64 + 64 = 128. The Router uses the TE metric by default for path selection process and the te metric matches the igp metric by default.

The section Config Parameters section covers "Bandwidth: 30". This reflects the bandwidth required by the tunnel to be established in kbps and it is taken from the (Global) pool of bandwidth from the interface.

The "setup priority" is listed here for quick reference.

Tunnel Affinity is covered later but is noted here as a 32 bit value if needed.

"Metric Type is "TE" by default as covered in the explanation of the basis for setup.

"Autoroute- enable" was the chosen method to route traffic into the tunnel. The "Lockdown-disabled" feature does not allow the tunnel to move from the current path and is disabled by default.

The "Load-share- disabled" can be used instead of specifying the bandwidth manually. It will load share the traffic between tunnels if more than one tunnel exist and has this feature enabled as well. There is further explanation of (bw-based) protocols later in the lesson. "auto-bw- disabled" field can be used in conjunction with manual bandwidth statements but is disabled by default.

The "Inlabel" displays the label as it applies to the data plane. A label that arrives from an peer is swapped with the "OutLabel" that may be imposed initially, as in this case. The "Outlabel" displays the label that is forwarded toward the route origination point and specifies the interface (Serial 0/0.12) and the label (16) that is used to accomplish this task.

The "RSVP signalling Info:" displays the SouRCe and DeSTination of the tunnel RSVP configured tunnel. The tunnel ID is "0" and the "Tun_Instance" will change anytime the tunnel is torn down.

The "RSVP Path Info:" list the address of the router and the "Explicit Route" of the hop-by-hop path from R1 to R4 through R2. "Record Route" will be covered later.
The Tspec , or Traffic specification is used to enforce policing on the traffic using the tunnel to the

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Tspec value specified in kbps. The Fspec or Flow specification is the value that will be used in RESerVation (RESV) messages by a receiver node to request QOS parameters.

The History of the tunnel can prove very useful in troubleshooting to establish patterns of success or failure. The "Tunnel:" is the time since the tunnel was created and how long since the last path change.

The "Current LSP:" output lists how long the current tunnel has been operational.

The "Prior LSP:" shows the path option in question and the "Removal Trigger" list what the potential problem or cause of the last LSP change or failure. The prior instance" is listed between the brackets after the path option.

```
R1#Sh ip rsvp int

interface     allocated  i/f max   flow max     sub max
F0/0          0          60K       60K          0
Se0/0.12      30K        60K       60K          0
Se0/0         0          0         0            0
Se0/0.13      0          60K       60K          0
```

The interface through R2 is reserving the bandwidth as 30K.

Control and Data Plane:

On R1

```
R1#Sh mpls ldp parameters

Protocol version: 1
No label generic region for downstream label distribution
Session hold time: 180 sec; keep alive interval: 60 sec
Discovery hello: holdtime: 15 sec; interval: 5 sec
Discovery targeted hello: holdtime: 90 sec; interval: 10 sec
Downstream on Demand max hop count: 255
Downstream on Demand Path Vector Limit: 255
LDP for targeted sessions
LDP initial/maximum backoff: 15/120 sec
LDP loop detection: off
```

R1 is performing "Downstream on Demand" learning of labels. This means that R1 will need to receive a label concerning a Downstream Network first before allocating a label for a given destination.

On R4

```
R4#Sh mpls traffic-eng tunnels int  | b InLabel |Out
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

InLabel : Serial0/0.42, implicit-null
OutLabel : -

R1 sends an RSVP "PATH message" that request for a label concerning the tunnel.
The request passes through R2 and eventually reaches R4. Once the request is received for the tunnel, it is assigned a label of "implicit null" and will send this label to R2. The "OutLabel" represents the action that will be taken in the event it receives the "InLabel" context. In this case, this is the termination point so there should be no Outlabel Context.

On R2

```
R2#Sh mpls traffic-eng tunnels int  | b InLabel |Out

  InLabel  : Serial0/0.21, 16
  OutLabel : Serial0/0.24, implicit-null

R2#Sh mpls traffic-eng tun brie | b TUNNEL

TUNNEL NAME              DESTINATION     UP IF     DOWN IF   STATE/PROT
R1_t0                    4.4.4.4         Se0/0.21  Se0/0.24  up/up
Displayed 0 (of 0) heads, 1 (of 1) midpoints, 0 (of 0) tails

R2#Sh mpls forwarding-table | b Local

Local  Outgoing    Prefix          Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id    switched   interface
16     Pop tag     1.1.1.1 0 [4]   0          Se0/0.24   point2point
```

Control Plane: R2 assignes a local tag of "16" for the tunnel when it receives the implicit null tag from R4. The tag could be different in your output.

Data Plane: R2 is waiting for label "16" on the InLabel interface s0/0.21 and forwarding the route with the implicit null tag (label "3") out of "OutLabel" int s0/0.24.

On R1

```
R1#Sh mpls traffic-eng tunnels int  | b InLabel |Out
```

InLabel : -
OutLabel : Serial0/0.12, 16

Control Plane: R1 learns the "16 label from R2.
Data Plane: Label "16" will be imposed if R1 receives a packet destined to use the tunnel.

```
R1#Trace 4.4.4.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Tracing the route to 4.4.4.4

  1 10.1.12.2 [MPLS: Label 16 Exp 0] 48 msec 28 msec 4 msec
  2 10.1.24.4 4 msec *  4 msec

R1#Sh ip cef exact-route 1.1.1.1 4.4.4.4

1.1.1.1          -> 4.4.4.4          : Tunnel0 (next hop 0.0.0.0)
```

R1 uses the slower path due to the TE configuration on the R1 tunnel head.

Task 3

Configure a second tunnel (tunnel "1") on R1 that uses the same path option configured in Task 2. Provide an additional path option 2 MEDIUM_LINK on both tunnels. This path option should use the path through R3.

On R1

```
R1(config)#interface Tunnel0
R1(config-if)#tunnel mpls traffic-eng path-option 2 explicit name MEDIUM_LINK
```

We add the second path option on tunnel 0

```
R1(config)#interface Tunnel1
R1(config-if)#ip unnumbered Loopback0
R1(config-if)#tunnel destination 4.4.4.4
R1(config-if)#tunnel mode mpls traffic-eng
R1(config-if)#tunnel mpls traffic-eng autoroute announce
R1(config-if)#tunnel mpls traffic-eng priority 4 4
R1(config-if)#tunnel mpls traffic-eng bandwidth  30
R1(config-if)#tunnel mpls traffic-eng path-option 1 explicit name SLOW_LINK
R1(config-if)#tunnel mpls traffic-eng path-option 2 explicit name MEDIUM_LINK
```

The second tunnel is created with similar options for simplicity. The path will not come up until at least one of the explicit paths is configured.

```
R1(config)#ip explicit-path name MEDIUM_LINK
R1(cfg-ip-expl-path)#next-address 10.1.13.3

Explicit Path name MEDIUM_LINK:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    1: next-address 10.1.13.3
R1(cfg-ip-expl-path)#next-address 10.1.34.4
Explicit Path name MEDIUM_LINK:
    1: next-address 10.1.13.3
    2: next-address 10.1.34.4

R1(cfg-ip-expl-path)#end
```

The tunnel comes up immediately. 2 paths have been created. We need to see what tunnel is being used and analyze the constraints.

```
R1#Sh mpls traffic-eng tunnels name R1_t0

Name: R1_t0                           (Tunnel0) Destination: 4.4.4.4
  Status:
    Admin: up        Oper: up     Path: valid       Signalling: connected

    path option 1, type explicit SLOW_LINK (Basis for Setup, path weight 128)
    path option 2, type explicit MEDIUM_LINK

  Config Parameters:
    Bandwidth: 30       kbps (Global)  Priority: 4  4   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 30       bw-based
    auto-bw: disabled

  InLabel  :  -
  OutLabel : Serial0/0.12, 16
  RSVP Signalling Info:
       Src 1.1.1.1, Dst 4.4.4.4, Tun_Id 0, Tun_Instance 26
    RSVP Path Info:
      My Address: 1.1.1.1
      Explicit Route: 10.1.12.2 10.1.24.4 4.4.4.4

---text omitted---
```

The "Tunnel 0" interface uses the 1st path option since the tunnel priority does not come into play in this scenario and the bandwidth is available.

```
R1#Sh mpls traffic-eng tunnels name-regexp ^R1_t1$
```

Yes….We can use regular expressions to match on names!!

```
Name: R1_t1                           (Tunnel1) Destination: 4.4.4.4

  Status:
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
  Admin: up        Oper: up     Path: valid       Signalling: connected

  path option 1, type explicit SLOW_LINK (Basis for Setup, path weight 128)
  path option 2, type explicit MEDIUM_LINK

 Config Parameters:
  Bandwidth: 30      kbps (Global)  Priority: 4  4   Affinity: 0x0/0xFFFF
  Metric Type: TE (default)
  AutoRoute:  enabled   LockDown: disabled  Loadshare: 30       bw-based
  auto-bw: disabled

 InLabel  :  -
 OutLabel : Serial0/0.12, 17
 RSVP Signalling Info:
     Src 1.1.1.1, Dst 4.4.4.4, Tun_Id 1, Tun_Instance 1
  RSVP Path Info:
    My Address: 1.1.1.1
    Explicit Route: 10.1.12.2 10.1.24.4 4.4.4.4

---text omitted---
```

"Tunnel 1" also uses the SLOW_LINK. There is enough bandwidth available and it is also the first option on this tunnel as well.

Task 4

Configure Tunnel "1 "to use 40 K of bandwidth and reoptomize the tunnels. Since RSVP bandwidth has been limited to 60K on every interface, view the changes that result in reoptimization of the tunnels.

On R1

```
R1(config)#interface tunnel 1
R1(config-if)#tunnel mpls traffic-eng bandwidth 40
```

We will view the effects of changing the bandwidth requirement of "Tunnel 1"

```
R1#Sh mpls traffic-eng tunnels brie

Signalling Summary:
   LSP Tunnels Process:          running
   RSVP Process:                 running
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Forwarding:                 enabled
  Periodic reoptimization:    every 3600 seconds, next in 1585 seconds
  Periodic auto-bw collection:  disabled
TUNNEL NAME       DESTINATION     UP IF     DOWN IF   STATE/PROT
R1_t0             4.4.4.4           -       Se0/0.12    up/up
R1_t1             4.4.4.4           -       Se0/0.13    up/up
Displayed 2 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
```

The traffic that is currently using the SLOW_LINK may not change until reoptimization takes place (3600 seconds). In ths case, a changing of the resources has taken place, which triggers TE path change due to the constraint. Disabling timer frequency or events will have no effect.

```
R1#Sh mpls traffic-eng tunnel role head source-id 1.1.1.1 name R1_t1

Name: R1_t1                           (Tunnel1) Destination: 4.4.4.4
  Status:
    Admin: up       Oper: up     Path: valid      Signalling: connected

    path option 2, type explicit MEDIUM_LINK (Basis for Setup, path weight 128)
    path option 1, type explicit SLOW_LINK

  Config Parameters:
    Bandwidth: 40       kbps (Global)  Priority: 4  4   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 40       bw-based
    auto-bw: disabled

  InLabel  :  -
  OutLabel : Serial0/0.13, 16
  RSVP Signalling Info:
       Src 1.1.1.1, Dst 4.4.4.4, Tun_Id 1, Tun_Instance 3
    RSVP Path Info:
      My Address: 1.1.1.1
      Explicit Route: 10.1.13.3 10.1.34.4 4.4.4.4
      Record Route:  NONE
      Tspec: ave rate=40 kbits, burst=1000 bytes, peak rate=40 kbits
    RSVP Resv Info:
      Record Route:  NONE
      Fspec: ave rate=40 kbits, burst=1000 bytes, peak rate=40 kbits
  History:
    Tunnel:
      Time since created: 5 minutes, 30 seconds
      Time since path change: 2 minutes, 6 seconds
    Current LSP:
      Uptime: 2 minutes, 6 seconds
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Selection: reoptimation
Prior LSP:
ID: path option 1 [1]
Removal Trigger: configuration changed
Last Error: PCALC:: Can't use link 0.0.0.0 on node 1.1.1.1

Notice how although the SLOW_LINK is the 1st option, the MEDIUM_PATH is chosen because the primary link only has a bandwidth of 60 kbps configured on the interface. The "Tunnel 0" interface has a requirement of 30 kbps. The Tunnel 1" interface has a requirement of 40 kbps and RSVP will reject the reservation on the primary path. The second path is unused and is chosen as the best for this tunnel.
Also, the Current LSP information displays that the new LSP was due to reoptimization. The reoptimization is now automatic when a change is performed on the tunnel and does not necessarily need to be manually triggered. In Cisco IOS, The lowered number tunnel will win if ALL constraints are tied except for bandwidth.

Task 5

Configure the tunnels such that when shutdown manually, Tunnel"1" will use the primary path no matter which order the tunnels come up. Even during reoptimization.

On R1

```
R1(config)#int tun 0
R1(config-if)#shut
R1(config-if)#int tun 1
R1(config-if)#shut

R1(config)#int tunnel 1
R1(config-if)#tunnel mpls traffic-eng priority 3 4
% Setup priority (3) may not be higher than hold priority (4)

R1(config-if)#tunnel mpls traffic-eng priority 3 3
```

The priority is used to accomplish the task correctly. The hold priority cannot be a higher number (inferior) than the setup priority. They are normally set to be identical but the hold priority can be set lower (better) than the setup priority. In some cases the traffic using the tunnel may not need a better priority to be established but once it is established it should not be taken down.

```
R1(config)#int tun 0
R1(config-if)#NO shut
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*Mar  1 00:48:26.939: %LINK-3-UPDOWN: Interface Tunnel0, changed state to up
*Mar  1 00:48:27.939: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed state to up

R1#Sh mpls traffic-eng tunnel brie | b TUNNEL

TUNNEL NAME           DESTINATION    UP IF     DOWN IF   STATE/PROT
R1_t0                 4.4.4.4        -         Se0/0.12  up/up
R1_t1                 4.4.4.4        -         unknown   admin-down
Displayed 2 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
```

We bring up tunnel "0" first and make sure it is established. It uses the primary path through R2.

```
R1(config)#int tun 1
R1(config-if)#NO shut
```

Verifying the Priority constraint:

On R1

```
R1#Sh mpls traffic-eng tunnels name R1_t1 | include path option

  path option 1, type explicit SLOW_LINK (Basis for Setup, path weight 128)
  path option 2, type explicit MEDIUM_LINK
    ID: path option 2 [16]
```

As we can see. Tunnel "1" gets established along the primary path. It bumps off Tunnel 0" in due to the better setup priority.

```
R1#Sh mpls traffic-eng tunnels name R1_t0 | include path option

  path option 2, type explicit MEDIUM_LINK (Basis for Setup, path weight 128)
  path option 1, type explicit SLOW_LINK
    ID: path option 1 [33]
```

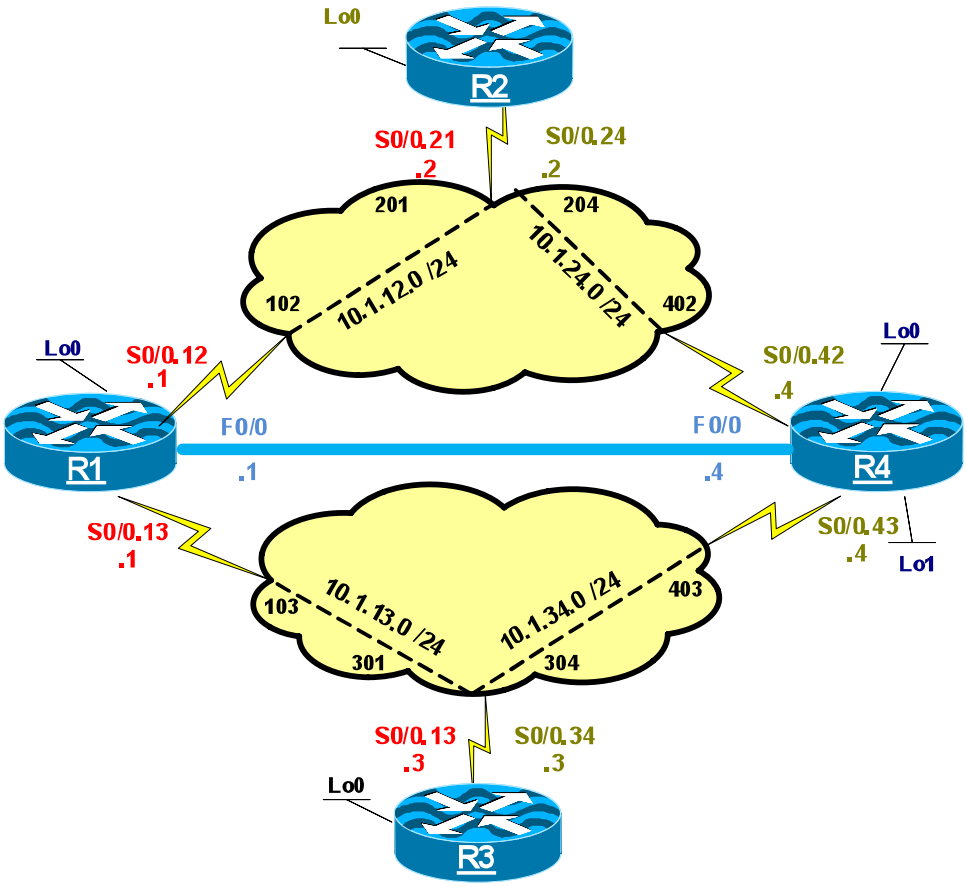
Tunnel "0" is setup across the MEDIUM_LINK path instead.

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 2 – Dynamic Tunnels



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12 S0/0.13 F0/0	1.1.1.1 /32 10.1.12.1 /24 10.1.13.1 /24 10.1.14.1 /24
R2	Lo0 S0/0.21 S0/0.23	2.2.2.2 /32 10.1.12.2 /24 10.1.24.2 /24
R3	Lo0 S0/0.31 S0/0.34	3.3.3.3/ 32 10.1.13.3 /24 10.1.34.3 /24
R4	Lo0 Lo1 F0/0 S0/0.42 S0/0.43	4.4.4.4 /32 44.4.4.4 /32 10.1.14.4 /24 10.1.24.4 /24 10.1.34.4 /24

Lab Setup:

The lab must have basic IGP and MPLS traffic Engineering configured to support tunnels. RSVP should use a bandwidth of 80 Kbps unless noted.

Task 1

Configure an OSPF with a process of "1" on all routers. The routers must be prepared to support MPLS-TE. Complete connectivity between the devices is required. The bandwidth for the R1-R3-R4 path should be set to 768K.

On R1

```
R1(config)#mpls traffic-eng tunnels

R1(config)#router ospf 1
R1(config)#network 0.0.0.0 0.0.0.0 area 0
R1(config-router)#mpls traffic-eng area 0
R1(config-router)#mpls traffic-eng router-id loopback 0

R1(config)#interface s0/0.12
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth 80

R1(config)#interface f0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth 80

R1(config)#interface s0/0.13
R1(config-if)#bandwidth 768
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth 80
```

On R2

```
R2(config)#mpls traffic-eng tunnels

R2(config)#router ospf 1
R2(config)#network 0.0.0.0 0.0.0.0 area 0
R2(config-router)#mpls traffic-eng area 0
R2(config-router)#mpls traffic-eng router-id loopback 0

R2(config)# interface s0/0.21
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth 80

R2(config-if)#interface s0/0.24
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth 80
```

On R3

```
R3(config)#mpls traffic-eng tunnels

R3(config)#router ospf 1
R3(config)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#mpls traffic-eng area 0
R3(config-router)#mpls traffic-eng router-id loopback 0

R3(config)# interface s0/0.31
R3(config-if)#bandwidth 768
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth 80

R3(config-if)#interface s0/0.34
R3(config-if)#bandwidth 768
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth 80
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4(config)#mpls traffic-eng tunnels

R4(config)#router ospf 1
R4(config)#network 0.0.0.0 0.0.0.0 area 0
R4(config-router)#mpls traffic-eng area 0
R4(config-router)#mpls traffic-eng router-id loopback 0

R4(config)#interface s0/0.42
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth 80

R4(config)# interface f0/0
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth 80

R4(config-if)#interface s0/0.43
R4(config-if)#bandwidth 768
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth 80
```

Verify the Configuration

On R1

```
R1#Sh ip route ospf | i O

O       2.2.2.2 [110/65] via 10.1.12.2, 00:00:09, Serial0/0.12
O       4.4.4.4 [110/11] via 10.1.14.4, 00:00:09, FastEthernet0/0
O       10.1.24.0 [110/74] via 10.1.14.4, 00:00:09, FastEthernet0/0
O       10.1.34.0 [110/140] via 10.1.14.4, 00:00:09, FastEthernet0/0
O       44.4.4.4 [110/11] via 10.1.14.4, 00:00:09, FastEthernet0/0

R1#Sh mpls int

Interface                  IP            Tunnel    Operational
FastEthernet0/0            No            Yes         Yes
Serial0/0.12               No            Yes         Yes
Serial0/0.13               No            Yes         Yes
```

On R2

```
R2#Sh ip route ospf | i O
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       1.1.1.1 [110/65] via 10.1.12.1, 00:02:07, Serial0/0.21
O       3.3.3.3 [110/195] via 10.1.24.4, 00:02:07, Serial0/0.24
O       4.4.4.4 [110/65] via 10.1.24.4, 00:02:07, Serial0/0.24
O       10.1.14.0 [110/74] via 10.1.24.4, 00:02:07, Serial0/0.24
O       10.1.13.0 [110/194] via 10.1.12.1, 00:02:07, Serial0/0.21
O       10.1.34.0 [110/194] via 10.1.24.4, 00:02:07, Serial0/0.24
O       44.4.4.4 [110/65] via 10.1.24.4, 00:02:07, Serial0/0.24

R2#Sh mpls int

Interface                  IP             Tunnel      Operational
Serial0/0.21               No             Yes         Yes
Serial0/0.24               No             Yes         Yes
```

The Tunnel Section is specifically for the TE configuration.

On R3

```
R3#Sh ip route ospf | i O

O       1.1.1.1 [110/131] via 10.1.13.1, 00:02:44, Serial0/0.31
O       2.2.2.2 [110/195] via 10.1.34.4, 00:02:44, Serial0/0.34
O       4.4.4.4 [110/131] via 10.1.34.4, 00:02:44, Serial0/0.34
O       10.1.14.0 [110/140] via 10.1.34.4, 00:02:44, Serial0/0.34
O       10.1.12.0 [110/194] via 10.1.13.1, 00:02:44, Serial0/0.31
O       10.1.24.0 [110/194] via 10.1.34.4, 00:02:44, Serial0/0.34
O       44.4.4.4 [110/131] via 10.1.34.4, 00:02:44, Serial0/0.34

R3#Sh mpls int

Interface                  IP             Tunnel      Operational
Serial0/0.31               No             Yes         Yes
Serial0/0.34               No             Yes         Yes
```

On R4

```
R4#Sh ip route ospf | i O

O       1.1.1.1 [110/11] via 10.1.14.1, 00:03:11, FastEthernet0/0
O       2.2.2.2 [110/65] via 10.1.24.2, 00:03:11, Serial0/0.42
O       3.3.3.3 [110/131] via 10.1.34.3, 00:03:11, Serial0/0.43
O       10.1.13.0 [110/140] via 10.1.14.1, 00:03:11, FastEthernet0/0
O       10.1.12.0 [110/74] via 10.1.14.1, 00:03:11, FastEthernet0/0

R4#Sh mpls int
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Interface                IP              Tunnel      Operational
FastEthernet0/0          No              Yes         Yes
Serial0/0.42             No              Yes         Yes
Serial0/0.43             No              Yes         Yes
```

Task 2

Create a tunnel from R1 to R4 (Tunnel0) that automatically builds according to the current cost of the links. The signalled bandwidth should be 50K. Tunnel Priority should be setup of 1 and hold of 1. All traffic from any source that terminates to 4.4.4.4 should use the tunnel.

On R1

```
R1(config)#interface Tunnel0
R1(config-if)#ip unnumbered Loopback0
R1(config-if)#tunnel destination 4.4.4.4
R1(config-if)#tunnel mode mpls traffic-eng
R1(config-if)#tunnel mpls traffic-eng autoroute announce
R1(config-if)#tunnel mpls traffic-eng priority 1 1
R1(config-if)#tunnel mpls traffic-eng bandwidth  50
R1(config-if)#tunnel mpls traffic-eng path-option 1 dynamic
```

The path option uses the "dynamic" keyword instead of static to accomplish the task.

Verify the dynamic tunnel:

```
R1#Show mpls traffic-eng tunnels up | s RSVP Path Info:

  RSVP Path Info:
    My Address: 10.1.14.1
    Explicit Route: 10.1.14.4 4.4.4.4
    Record Route:  NONE
    Tspec: ave rate=50 kbits, burst=1000 bytes, peak rate=50 kbits
```

The path direct to R4 is chosen as the best path with the necessary bandwidth available.

```
R1#Sh mpls traffic-eng link-management admission-control f0/0

System Information::
  Tunnels Count:       1
  Tunnels Selected:    1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
TUNNEL ID      UP IF    DOWN IF  PRIORITY    STATE          BW (kbps)
1.1.1.1 0_1    -        Fa0/0    1/1         Resv Admitted  50          RG

R1#Sh ip rsvp int f0/0

interface       allocated      i/f max      flow max     sub max
Fa0/0           50K            80K          80K          0
```

30K more can be allocated on this interface.

Task 2

Configure a second tunnel called Tunnel "10 "to use 40 K of bandwidth. The priority should be identical to tunnel 0. Traffic destined for 44.4.4.4 should use this tunnel instead of tunnel "0". Policy-Based routing is not permitted.

On R1

```
R1(config)#interface Tunnel10
R1(config-if)#ip unnumbered Loopback0
R1(config-if)#tunnel destination 4.4.4.4
R1(config-if)#tunnel mode mpls traffic-eng
R1(config-if)#tunnel mpls traffic-eng priority 1 1
R1(config-if)#tunnel mpls traffic-eng bandwidth  40
R1(config-if)#tunnel mpls traffic-eng path-option 1 dynamic
R1(config-if)#exit
```

The bandwidth constraint has been raised to 40 to force a calculation to be made on which tunnel gets to use the best path. The primary path has now been overutilized by 10K.

```
R1(config)#ip route 44.4.4.4 255.255.255.255 tunnel 10
```

This option can be used instead of routing traffic with autoroute announce

Verify Path Info:

```
R1#Sh ip cef exact-route 1.1.1.1 44.4.4.4

1.1.1.1          -> 44.4.4.4        : Tunnel10 (attached)

R1#Sh ip cef exact-route 1.1.1.1 4.4.4.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
1.1.1.1          -> 4.4.4.4          : Tunnel0 (next hop 0.0.0.0)
```

The "autoroute announce" option on "Tunnel 0" allows ALL traffic that is destined for 4.4.4.4 to use tunnel 0. The static overrides this constraint by choosing to use "Tunnel 10" for routes to the 44.4.4.4 destination prefix.

```
R1#Sh mpls traffic-eng tunnels name R1_t10 | section RSVP Path Info:

    RSVP Path Info:
      My Address: 1.1.1.1
      Explicit Route: 10.1.12.2 10.1.24.4 4.4.4.4
      Record Route:  NONE
      Tspec: ave rate=40 kbits, burst=1000 bytes, peak rate=40 kbits
```

The path through R2 was chosen due to the fact that Tunnel 0 was already established with a tunnel priority of "1" for setup and "1" for hold. "Tunnel 10" attempts to signal but the bandwidth constraint is not sufficient for the path to R4 directly. The tunnel reroutes all traffic through R2 since it is the next best path with a cost of 64.

Task 3

Tunnel 10 should take the path through R1, R3 and R4. Priority, Policy based Routing explicit paths or modifying OSPF cost is not permitted.

Viewing the current tunnel state:

```
R1#Sh mpls traffic-eng link-management advertisements | sec Link ID::  1

Link ID:: 1
   Link IP Address:     10.1.12.1
   IGP Neighbor:        ID 2.2.2.2, IP 10.1.12.2
   TE metric:           64
   IGP metric:          64
   Physical Bandwidth:  1544 kbits/sec
   Res. Global BW:      80 kbits/sec
   Res. Sub BW:         0 kbits/sec
   Downstream::

---text omitted---

R1#Sh mpls traffic-eng tunnels name R1_t10 | sec  Config Parameters:
```

2 spaces

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
 Config Parameters:

    Bandwidth: 40       kbps (Global)  Priority: 1  1   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute:  disabled  LockDown: disabled  Loadshare: 40       bw-based
    auto-bw: disabled
```

TE metric is another cost that can be used instead of the IGP metric. By default, the tunnel TE metrics also equals the IGP metric. (TE = IGP).

```
R1#Show mpls traffic-eng link-management bandwidth-allocation | inc
Admin. Weight:

      Admin. Weight:          10 (IGP)
      Admin. Weight:          64 (IGP)
      Admin. Weight:          130 (IGP)

R1#Show mpls traffic-eng link-managemen interfaces s0/0.13 | incl Admin. Weight:

      Admin. Weight:          130 (IGP)
```

The cost of all 3 links from R1 to R4 is as follows: F0/0 = *10*, S0/0.12 = *64* and S0/0.13 = *130*.

On R3

```
R3(config)#interface s0/0.34
R3(config-subif)#mpls traffic-eng administrative-weight 1
```

On R1

```
R1(config)#interface s0/0.13
R1(config-subif)#mpls traffic-eng administrative-weight 1

R1(config)# interface tunnel 0
R1(config-if)#tunnel mpls traffic-eng path-selection metric igp
R1(config-if)#end

R1#mpls traffic-eng reoptimize

R1#sh mpls traffic-eng tunnels name R1_t0 | sec  RSVP Path Info:

    RSVP Path Info:
      My Address: 10.1.14.1
      Explicit Route: 10.1.14.4 4.4.4.4
      Record Route:  NONE
      Tspec: ave rate=50 kbits, burst=1000 bytes, peak rate=50 kbits
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Sh mpls traffic-eng tunnels name R1_t10 | be RSVP Path Info:

  RSVP Path Info:
       My Address: 1.1.1.1
       Explicit Route: 10.1.13.3 10.1.34.4 4.4.4.4

R1#Show mpls traffic-eng link-managemen interfaces s0/0.13 | incl Admin.
Weight:

Admin. Weight:          1 (configured)

R1#Sh mpls traffic-eng tunnels name R1_t10 | s Status

  Status:
    Admin: up          Oper: up     Path: valid       Signalling: connected
    path option 1, type dynamic (Basis for Setup, path weight 2)
```

Tunnel "0" uses the path through the lowest IGP cost which still uses the path to R4 directly. Tunnel"10" uses the path that reflects the lowest TE metric of "2".

Task 4

Tunnel 0 should take the path through R2 to R4. Priority, Policy based Routing, explicit paths modifying OSPF cost or TE metric is not permitted.

Viewing the current tunnel paths:

```
R1#Trace 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4

  1 10.1.14.4 8 msec *  4 msec

R1#Trace 44.4.4.4

Type escape sequence to abort.
Tracing the route to 44.4.4.4

  1 10.1.13.3 [MPLS: Label 16 Exp 0] 52 msec 4 msec 24 msec
  2 10.1.34.4 24 msec *  8 msec
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Currently, the path for 4.4.4.4 is across the f0/0 interface directly to R4. This tunnel routes ALL traffic behind it to Tunnel 0.

On R1

```
R1(config-if)#int s0/0.12
R1(config-if)#mpls traffic-eng attribute-flags 0xAD

R1(config)#int Tunnel 0
R1(config-if)#tunnel mpls traffic-eng affinity 0x1 mask 0x3
```

On R2

```
R2(config-if)#int s0/0.21
R2(config-if)#mpls traffic-eng attribute-flags 0xAD

R2(config-if)#int s0/0.24
R2(config-if)#mpls traffic-eng attribute-flags 0xAD
```

Choosing the hexadecimal notation must be consistent if used. The other method is to provide the binary equivalent of 0x00110001 (0x31). This also must be consistent along the path. The tunnel should tear down when modifying the parameters. A reoptimization should not be necessary.

1) Only a "1" and a "1" will yield a "1" in the output when the Attribute Flag or Affinity is ANDED with the mask.

2) Affinity + Mask = Attribute flag + Mask

Tunnel 0

Affinity	0x00000001 (0x1)
Mask	0x00000011 (0x3)
Result	00000001

Path 1		Path 2	
Attribute flag	0x10110110 (0xB6)	Attribute flag	0x10101101 (0xAD)
Mask	0x00000011 (0x3)	Mask	0x00000011 (0x3)
Result	00000010 (no match)	Result	00000001 (match)

Path 1 was chosen to show a comparison of values when computing Affinity as a constraint. PATH 1 should not be used since it does not match the Result from the Tunnel Affinity and Mask. Only Path 2 matches and will be considered in path selection.

Verifying the tunnel path:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Sh mpls traffic-eng tunnels name R1_t0

Name: R1_t0                           (Tunnel0) Destination: 4.4.4.4
  Status:
    Admin: up       Oper: up     Path: valid      Signalling: connected

    path option 1, type dynamic (Basis for Setup, path weight 128)

  Config Parameters:
    Bandwidth: 50       kbps (Global)  Priority: 1  1   Affinity: 0x1/0x3
    Metric Type: IGP (interface)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 50       bw-based
    auto-bw: disabled

  InLabel  :  -
  OutLabel : Serial0/0.12, 16
  RSVP Signalling Info:
       Src 1.1.1.1, Dst 4.4.4.4, Tun_Id 0, Tun_Instance 19
    RSVP Path Info:
      My Address: 1.1.1.1
      Explicit Route: 10.1.12.2 10.1.24.4 4.4.4.4
      Record Route:  NONE
      Tspec: ave rate=50 kbits, burst=1000 bytes, peak rate=50 kbits
    RSVP Resv Info:
      Record Route:  NONE
      Fspec: ave rate=50 kbits, burst=1000 bytes, peak rate=50 kbits
  History:
    Tunnel:
      Time since created: 5 minutes, 40 seconds
      Time since path change: 2 minutes, 23 seconds
    Current LSP:
      Uptime: 2 minutes, 23 seconds
    Prior LSP:
      ID: path option 1 [12]
      Removal Trigger: configuration changed
```

The tunnel is restored and shown to have switched to the path where the values match. The Prior LSP shows that the configuration caused the switch to occur.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Allow Tunnel "0" to record the bandwidth every 30 seconds. Configure Tunnel"0" to have a tunnel bandwidth greater than "0" as a configuration under the tunnel interface without explicitly configuring it. The configuration should remain for at least 5 minutes. Tunnel 10 should be limited to no more than 90K of bandwidth at any time.

```
R1#Sh mpls traffic-eng tunnels name R1_t0 | sec Config Parameters:

  Config Parameters:
    Bandwidth: 50       kbps (Global)  Priority: 1  1   Affinity: 0x1/0x3
    Metric Type: IGP (interface)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 50       bw-based
    auto-bw: disabled
```

The auto-bandwidth feature is currently disabled and is the focus of this task.

On R1

```
R1(config)#mpls traffic-eng auto-bw timers frequency 30

R1(config)#int tunnel 0
R1(config-if)#tunnel mpls traffic-eng auto-bw frequency 300
```

After 300 seconds,

```
R1(config)#int tunnel 10
R1(config-if)#tunnel mpls traffic-eng auto-bw max-bw 90
```

Auto bandwidth can be configured with the following parameters:

"mpls traffic-eng auto-bw frequency" – This command when configured globally will record the bandwidth on ALL TE interfaces in a prescribed amount of time in seconds.

"tunnel mpls traffic-eng auto-bw"- This command changes optional parameters for auto-bw calculation. The default is 86400 seconds (1 day).

***collect-bw* – only states what the bandwidth would be when a change takes place.**

***frequency*- changes the time when the bandwidth is used from the recorded data on the TE interface. The router will automatically configure the requested bandwidth under the apprpriate tunnel.**

***max-bw* – maximizes how much bandwidth can be allocated for a tunnel when the frequency timer expires.**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

min-bw **- expresses the minimum bandwidth required for allocation to the tunnel when the frequency timer expires.**

Verify the configuration:

```
R1#Sh mpls traffic-eng tunnels name R1_t0 | sec Config Parameters:

  Config Parameters:
    Bandwidth: 50       kbps (Global)  Priority: 1  1   Affinity: 0x1/0x3
    Metric Type: IGP (interface)
    AutoRoute:  enabled   LockDown: disabled  Loadshare: 50       bw-based
    auto-bw: (300/283) 0  Bandwidth Requested: 50
```

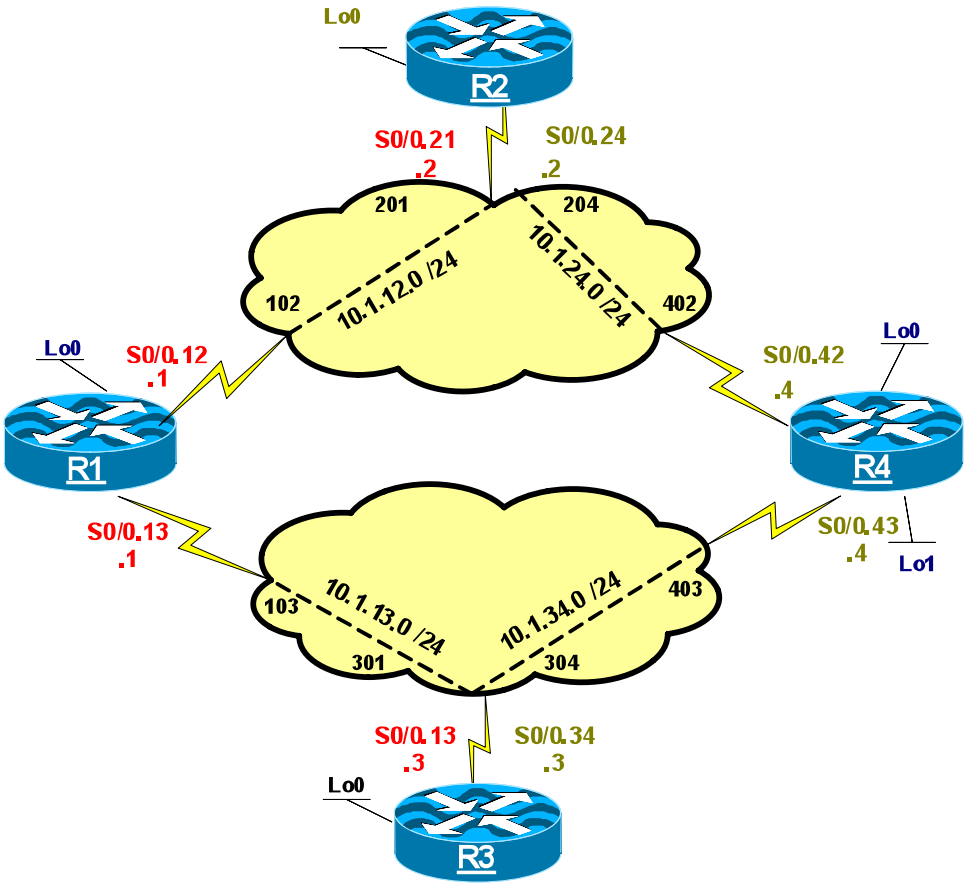
The 300 is the mpls frequency interval. The 285 represents the time left before a Bandwidth Calculation is retrived from the buffer. The first "0" is the current bandwidth that has been measured. The "Bandwidth Requested: 50" will post a number when the frequency timer expires and resets.

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 3 – Load Balancing



©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12 S0/0.13	1.1.1.1 /32 10.1.12.1 /24 10.1.13.1 /24
R2	Lo0 S0/0.21 S0/0.23	2.2.2.2 /32 10.1.12.2 /24 10.1.24.2 /24
R3	Lo0 S0/0.31 S0/0.34	3.3.3.3/ 32 10.1.13.3 /24 10.1.34.3 /24
R4	Lo0 Lo44 S0/0.42 S0/0.43	4.4.4.4 /32 44.44.44.44/32 10.1.24.4 /24 10.1.34.4 /24

Lab Setup:

All Frame-Relay links are connected in a point-to-point manner.

Task 1

Configure an OSPF with a process of "1" on all routers. The routers must be prepared to support MPLS-TE. Complete connectivity between the devices is required. All RSVP interfaces should be configured to allow 30K worth of traffic on ANY path.

On R1

```
R1(config)#mpls traffic-eng tunnels

R1(config)#router ospf 1
R1(config)#network 0.0.0.0 0.0.0.0 area 0
R1(config-router)#mpls traffic-eng area 0
R1(config-router)#mpls traffic-eng router-id loopback 0

R1(config)#interface s0/0.12
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth 30

R1(config-if)#interface s0/0.13
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth 30
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config)#mpls traffic-eng tunnels

R2(config)#router ospf 1
R2(config-router)#network 0.0.0.0 0.0.0.0 area 0
R2(config-router)#mpls traffic-eng area 0
R2(config-router)#mpls traffic-eng router-id loopback 0

R2(config)#interface s0/0.21
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth 30

R2(config-if)#interface s0/0.24
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth 30
```

On R3

```
R3(config)#mpls traffic-eng tunnels

R3(config)#router ospf 1
R3(config-router)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#mpls traffic-eng area 0
R3(config-router)#mpls traffic-eng router-id loopback 0

R3(config)#interface s0/0.31
R3(config-if)#ip ospf cost 63
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth 30

R3(config)#interface s0/0.34
R3(config-if)#ip ospf cost 63 (This will ensure that R1-R3-R4 is the best path)
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth 30
```

On R4

```
R4(config)#mpls traffic-eng tunnels

R4(config)#router ospf 1

R4(config-router)#network 0.0.0.0 0.0.0.0 area 0
R4(config-router)#mpls traffic-eng area 0
R4(config-router)#mpls traffic-eng router-id loopback 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config)#interface s0/0.42
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth 30

R4(config)#interface s0/0.43
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth 30
```

Verify the Configuration

On R1

```
R1#Sh ip route ospf | i O

O       2.2.2.2 [110/65] via 10.1.12.2, 00:00:09, Serial0/0.12
O       3.3.3.3 [110/65] via 10.1.13.3, 00:00:09, Serial0/0.13
O       4.4.4.4 [110/128] via 10.1.13.3, 00:00:09, Serial0/0.13
O       10.1.24.0 [110/128] via 10.1.12.2, 00:00:09, Serial0/0.12
O       10.1.34.0 [110/127] via 10.1.13.3, 00:00:09, Serial0/0.13
O       44.4.4.4 [110/128] via 10.1.13.3, 00:00:09, Serial0/0.13

R1#Sh mpls int

Interface                  IP            Tunnel    Operational
Serial0/0.12               No            Yes       Yes
Serial0/0.13               No            Yes       Yes
```

On R2

```
R2#Sh ip route ospf | i O

O       1.1.1.1 [110/65] via 10.1.12.1, 00:02:07, Serial0/0.21
O       3.3.3.3 [110/129] via 10.1.24.4, 00:02:07, Serial0/0.24
O       4.4.4.4 [110/65] via 10.1.24.4, 00:02:07, Serial0/0.24
O       10.1.13.0 [110/128] via 10.1.12.1, 00:02:07, Serial0/0.21
O       10.1.34.0 [110/128] via 10.1.24.4, 00:02:07, Serial0/0.24
O       44.4.4.4 [110/65] via 10.1.24.4, 00:02:07, Serial0/0.24

R2#Sh mpls int

Interface                  IP            Tunnel    Operational
Serial0/0.21               No            Yes       Yes
Serial0/0.24               No            Yes       Yes
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh ip route ospf | i O

O       1.1.1.1 [110/65] via 10.1.13.1, 00:02:44, Serial0/0.31
O       2.2.2.2 [110/128] via 10.1.34.4, 00:02:44, Serial0/0.34
O       4.4.4.4 [110/64] via 10.1.34.4, 00:02:44, Serial0/0.34
O       10.1.12.0 [110/128] via 10.1.13.1, 00:02:44, Serial0/0.31
O       10.1.24.0 [110/127] via 10.1.34.4, 00:02:44, Serial0/0.34
O       44.4.4.4 [110/64] via 10.1.34.4, 00:02:44, Serial0/0.34

R3#Sh mpls int

Interface                  IP             Tunnel      Operational
Serial0/0.31               No             Yes         Yes
Serial0/0.34               No             Yes         Yes
```

On R4

```
R4#Sh ip route ospf | i O

O       1.1.1.1 [110/128] via 10.1.34.3, 00:03:11, Serial0/0.43
O       2.2.2.2 [110/65] via 10.1.24.2, 00:03:11, Serial0/0.42
O       3.3.3.3 [110/65] via 10.1.34.3, 00:03:11, Serial0/0.43
O       10.1.13.0 [110/127] via 10.1.34.3, 00:03:11, Serial0/0.43
O       10.1.12.0 [110/128] via 10.1.24.2, 00:03:11, Serial0/0.42

R4#Sh mpls int

Interface                  IP             Tunnel      Operational
Serial0/0.42               No             Yes         Yes
Serial0/0.43               No             Yes         Yes
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 2

Create two new dynamic tunnels from R1 to R4 (4.4.4.4). The tunnels should have the attributes listed in the table below and should use autoroute announce. Tunnel "0" should be created first and should be torn down as a result of the second tunnel being built, however Tunnel 0 should not participate in reoptimization if the path through R3 becomes available.

Tunnel	Bandwidth	Priority (Setup/Hold)
Tunnel 0	30	7/7
Tunnel 1	30	6/6

Configuring the tunnels:

On R1

```
R1(config)#interface Tunnel0
R1(config-if)#ip unnumbered Loopback0
R1(config-if)#tunnel destination 4.4.4.4
R1(config-if)#tunnel mode mpls traffic-eng
R1(config-if)#tunnel mpls traffic-eng autoroute announce
R1(config-if)#tunnel mpls traffic-eng priority 7 7
R1(config-if)#tunnel mpls traffic-eng bandwidth  30
R1(config-if)#tunnel mpls traffic-eng path-option 1 dynamic lockdown

R1(config-if)#interface Tunnel1
R1(config-if)#ip unnumbered Loopback0
R1(config-if)#tunnel destination 4.4.4.4
R1(config-if)#tunnel mode mpls traffic-eng
R1(config-if)#tunnel mpls traffic-eng autoroute announce
R1(config-if)#tunnel mpls traffic-eng priority 6 6
R1(config-if)#tunnel mpls traffic-eng bandwidth  30
R1(config-if)#tunnel mpls traffic-eng path-option 1 dynamic
```

The lockdown option is used when a tunnel establishes connectivity and does not participate when an "mpls traffic-eng reoptimize" command is issued.

Verifying the Tunnel Setup:

```
R1#Sh mpls traffic-eng tunnels
```

Name: R1_t0 (Tunnel0) Destination: 4.4.4.4

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Status:
Admin: up Oper: up Path: valid Signalling: connected

path option 1,(LOCKDOWN) type dynamic (Basis for Setup, path weight 128)

---text omitted---
RSVP Path Info:
My Address: 1.1.1.1
Explicit Route: 10.1.12.2 10.1.24.4 4.4.4.4

Name: R1_t1 (Tunnel1) Destination: 4.4.4.4
Status:
Admin: up Oper: up Path: valid Signalling: connected

---text omitted---

RSVP Path Info:
My Address: 1.1.1.1
Explicit Route: 10.1.13.3 10.1.34.4 4.4.4.4

Tunnel "1" has a better priority and gets first path choice which is R1-R3-R4.
Tunnel "0 has taken the path through R2.

```
R1(config)#int tunnel 1
R1(config-if)#shut

R1#mpls traffic-eng reoptimize
```

Name: R1_t0 (Tunnel0) Destination: 4.4.4.4
Status:
Admin: up Oper: up Path: valid Signalling: connected

path option 1,(**LOCKDOWN**) type dynamic (Basis for Setup, path weight 128)

Explicit Route: 10.1.12.2 10.1.24.4 4.4.4.4

After shutting down Tunnel"1", the bandwidth becomes available on the best path.
When forcing the reoptimize feature, the bandwidth does not get reallocated for Tunnel"0" due to the "(LOCKDOWN)" status.

```
R2(config)#int s0/0.24
R2(config-subif)#shut
```

Here, we are simulating a failure by shutting down the interface somewhere in the path of Tunnel "0".

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Name: R1_t0                          (Tunnel0) Destination: 4.4.4.4
  Status:
    Admin: up        Oper: up     Path: valid       Signalling: connected

    path option 1,(LOCKDOWN) type dynamic (Basis for Setup, path weight 128)

  Config Parameters:
    Bandwidth: 30       kbps (Global)  Priority: 7  7   Affinity: 0x0/0xFFFF
    Metric Type: TE (default)
    AutoRoute:  enabled   LockDown: enabled   Loadshare: 30       bw-based
    auto-bw: disabled

  InLabel  :  -
  OutLabel : Serial0/0.13, 16
  RSVP Signalling Info:
       Src 1.1.1.1, Dst 4.4.4.4, Tun_Id 0, Tun_Instance 6
    RSVP Path Info:
      My Address: 1.1.1.1
      Explicit Route: 10.1.13.3 10.1.34.4 4.4.4.4

---text omitted---

History:
    Tunnel:
      Time since created: 7 minutes, 8 seconds
      Time since path change: 29 seconds
    Current LSP:
      Uptime: 29 seconds
      Selection: reoptimation
    Prior LSP:
      ID: path option 1 [5]
      Removal Trigger: path error
```

The path of the tunnel has changed to be what it would not reoptimize earlier. If a true failure takes place along the path, the tunnel will attempt to find a better path.

Task 3

A tunnel must be created from R1 to R4. Ensure that 2 times as many packets go down the path through R3 compared to 1 packet through R2 . The route should be recorded in MPLS-TE.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Viewing the current tunnel paths:

Be sure to "no shut" tunnel "1" and the interface on R2 before continuing to the next lab.

On R1

```
R1#Sh mpls traffic-eng autoroute 4.4.4.4

MPLS TE autorouting enabled
  destination 44.4.4.4 has 2 tunnels
    Tunnel0 (load balancing metric 66666666, nexthop 4.4.4.4)
    Tunnel1 (load balancing metric 66666666, nexthop 4.4.4.4)
```

The load balancing is calculated based on the tunnel bandwidth in kbp. (1,000,000 / 30 = 33,333.333). When this total is multiplied by itself, it equals the load sharing based on the output which = 66666666. Both paths have equal cost between them based on.

```
R1#Sh ip route ospf | sec 44.4.4.4

O        4.4.4.4 [110/128] via 0.0.0.0, 00:02:06, Tunnel0
                 [110/128] via 0.0.0.0, 00:02:06, Tunnel1
```

The cost to the 4.4.4.4 network prefix is currently 128. Cost = (64) between R1 and R2 + (64) Cost between R2 and R4 + Loopback = 1 (64+63+1 = 128)

```
R1(config)#int tun 0
R1(config-if)#tunnel mpls traffic-eng autoroute metric relative -5

R1#Sh mpls traffic-eng autoroute

MPLS TE autorouting enabled
  destination 4.4.4.4 has 2 tunnels
    Tunnel0 (load balancing metric 66666666, nexthop 4.4.4.4, relative metric -5)
    Tunnel1 (load balancing metric 66666666, nexthop 4.4.4.4)

R1#Sh ip route 44.4.4.4

Routing entry for 44.4.4.4/32
  Known via "ospf 1", distance 110, metric 123, type intra area
  Routing Descriptor Blocks:
  * directly connected, via Tunnel0
      Route metric is 123, traffic share count is 1
```

When the packet hits the end of the tunnel, 5 metric units are subtracted before counting each hop after. (64+63 – 5 = 123). In addition to the "*relative*" option, the "*absoute*" option will make the cost for everything after the tunnel to be the same no matter how many hops it takes to reach the destination

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

through the tunnel. This option only works when ISIS is used as the IGP.

```
R1(config)#int tun 0
R1(config-if)#tunnel mpls traffic-eng load-share 30

R1(config)#int tun 1
R1(config-if)#tunnel mpls traffic-eng load-share 60

R1#Sh mpls traffic-eng autoroute
```

MPLS TE autorouting enabled
  destination 4.4.4.4 has 2 tunnels
    Tunnel1 (**load balancing metric 33333333**, nexthop 4.4.4.4)
    Tunnel0 (**load balancing metric 66666666**, nexthop 4.4.4.4, relative metric -5

The load share can use any set of numbers to display the load balancing requested. 2 to 1 can also be represented by "200" and "100" respectively. The numbers may change but the load balancing is the same. The only difference is when the autoroute begins to load balance in relation to the bandwidth used.

Viewing the Tunnel prior to Route Record:

On R1

```
R1#Sh mpls traffic-eng tunnels name R1_t0 | sec RSVP Path Info:

    RSVP Path Info:
      My Address: 1.1.1.1
      Explicit Route: 10.1.12.2 10.1.24.4 4.4.4.4
      Record Route:  NONE
      Tspec: ave rate=30 kbits, burst=1000 bytes, peak rate=30 kbits
```

On Tunnel 1:

```
R1(config)#interface tunnel 1
R1(config-if)#tunnel mpls traffic-eng record-route
R1(config-if)#end

R1#mpls traffic-enf reoptimize
```

Viewing the Tunnel after Route Record enabled:

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4#Sh mpls traffic-eng tunnels name R1_t0 | sec RSVP Path Info:

    RSVP Path Info:
      My Address: 4.4.4.4
      Explicit Route:  NONE
      Record Route:  NONE
      Tspec: ave rate=30 kbits, burst=1000 bytes, peak rate=30 kbits

R4#Sh mpls traffic-eng tunnels name R1_t1 | sec RSVP Path Info:

    RSVP Path Info:
      My Address: 4.4.4.4
      Explicit Route:  NONE
      Record Route: 10.1.34.3 10.1.13.1
      Tspec: ave rate=30 kbits, burst=1000 bytes, peak rate=30 kbits
```

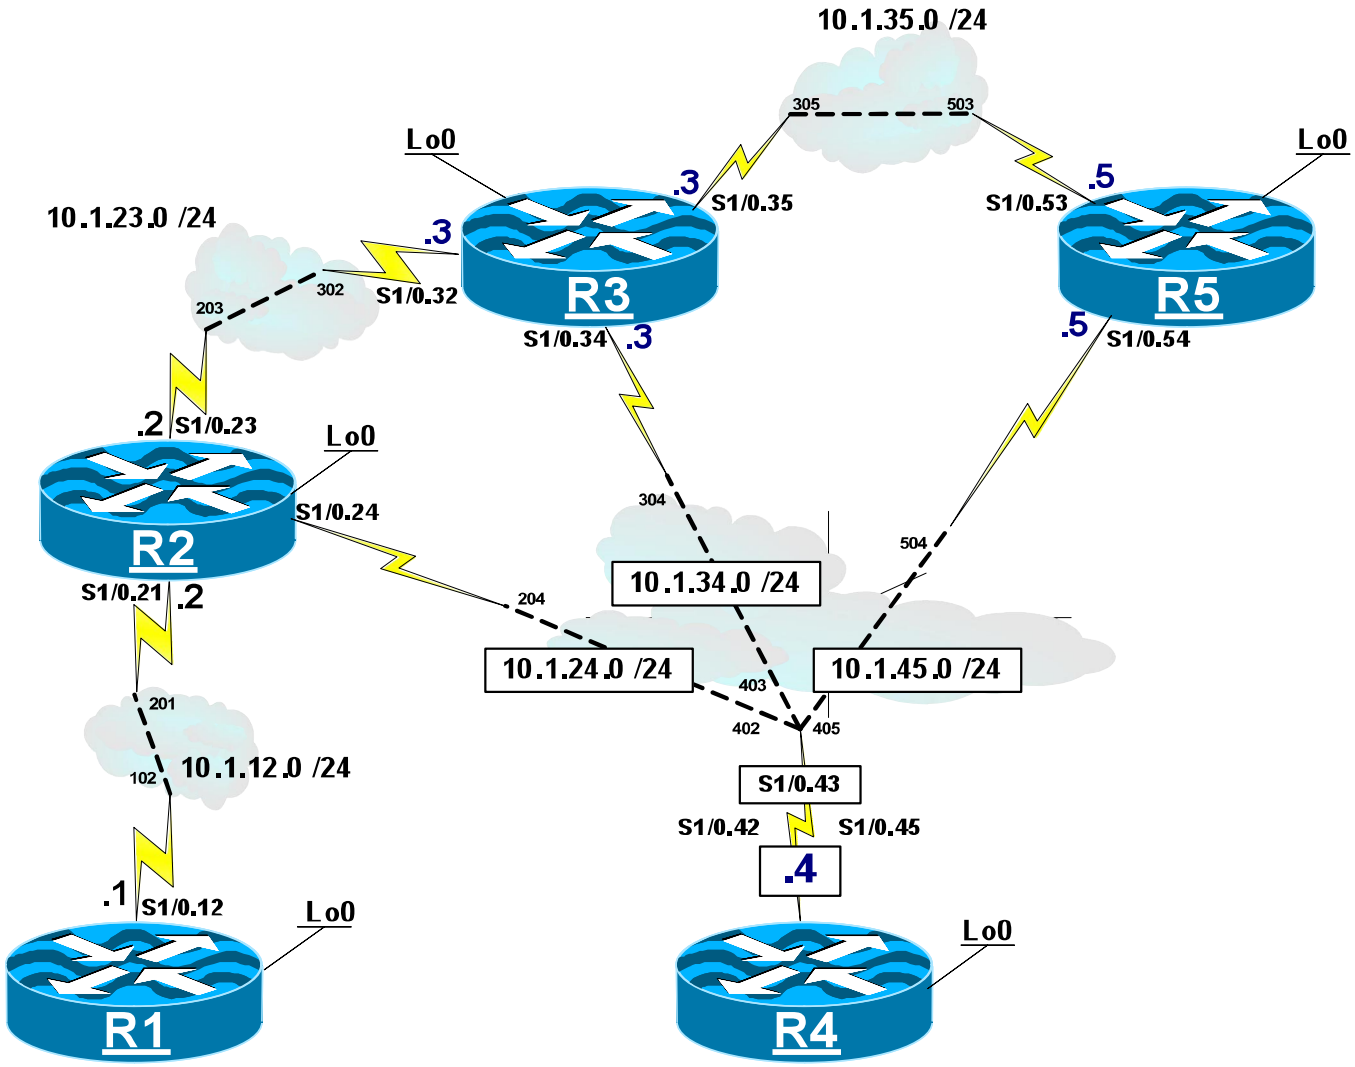
The Route Record option was only used on Tunnel "1". The route has been recorded and displayed on R4 as being valid.

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 4 –Fast reroute backup tunnels



IP Addressing:

Router	Interface	IP address
R1	Lo0 S1/0.12	1.1.1.1 /32 10.1.12.1 /24
R2	Lo0 S1/0.21 S1/0.23 S1/0.24	2.2.2.2 /32 10.1.12.2 /24 10.1.23.2/ 24 10.1.24.2 /24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

R3	Lo0 S1/0.32 S1/0.34 S1/0.35	3.3.3.3/ 32 10.1.23.3 /24 10.1.34.3. 24 10.1.35.3 /24
R4	Lo0 S1/0.42 S1/0.43 S1/0.45	4.4.4.4 /32 10.1.24.4 /24 10.1.34.4 /24 10.1.45.4 /24
R5	Lo0 S1/0.53 S1/0.54	5.5.5.5 /32 10.1.35.5 /24 10.1.45.5/ 24

Lab Setup:

All Frame-Relay links are connected in a point-to-point manner. The labels ranges should be manipulated for clarity.

Task 1

Configure an OSPF with a process of "1" on all routers. The routers must be prepared to support MPLS-TE. Complete connectivity between the devices is required. All RSVP interfaces should be configured to allow default bandwidth utilization.

On R1

```
R1(config)#mpls traffic-eng tunnels
R1(config)#mpls label range 100 199

R1(config)#router ospf 1
R1(config)#network 0.0.0.0 0.0.0.0 area 0
R1(config-router)#mpls traffic-eng area 0
R1(config-router)#mpls traffic-eng router-id loopback 0

R1(config)#interface S1/0.12
R1(config-if)#mpls traffic-eng tunnels
R1(config-if)#ip rsvp bandwidth
```

On R2

```
R2(config)#mpls traffic-eng tunnels
R2(config)#mpls label range 200 299

R2(config)#router ospf 1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#network 0.0.0.0 0.0.0.0 area 0
R2(config-router)#mpls traffic-eng area 0
R2(config-router)#mpls traffic-eng router-id loopback 0

R2(config)#interface S1/0.21
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth

R2(config)#interface S1/0.23
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth

R2(config)#interface S1/0.24
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth
```

On R3

```
R3(config)#mpls traffic-eng tunnels
R3(config)#mpls label range 300 399

R3(config)#router ospf 1
R3(config-router)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#mpls traffic-eng area 0
R3(config-router)#mpls traffic-eng router-id loopback 0

R3(config)#interface S1/0.32
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth

R3(config)#interface S1/0.34
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth

R3(config)#interface S1/0.35
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth
```

On R4

```
R4(config)#mpls traffic-eng tunnels
R4(config)#mpls label range 400 499

R4(config)#router ospf 1
R4(config-router)#network 0.0.0.0 0.0.0.0 area 0
R4(config-router)#mpls traffic-eng area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R4(config-router)#mpls traffic-eng router-id loopback 0

R4(config)#interface S1/0.42
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth

R4(config)#interface S1/0.43
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth

R4(config)#interface S1/0.45
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth
```

On R5

```
R5(config)#mpls traffic-eng tunnels
R5(config)#mpls label range 500 599

R5(config)#router ospf 1
R5(config-router)#network 0.0.0.0 0.0.0.0 area 0
R5(config-router)#mpls traffic-eng area 0
R5(config-router)#mpls traffic-eng router-id loopback 0

R5(config)#interface S1/0.53
R5(config-if)#mpls traffic-eng tunnels
R5(config-if)#ip rsvp bandwidth

R5(config)#interface S1/0.54
R5(config-if)#mpls traffic-eng tunnels
R5(config-if)#ip rsvp bandwidth
```

Verify the Configuration

On R1

```
R1#Sh ip route ospf | i O

O       2.2.2.2 [110/65] via 10.1.12.2, 00:00:53, Serial1/0.12
O       3.3.3.3 [110/129] via 10.1.12.2, 00:00:53, Serial1/0.12
O       4.4.4.4 [110/129] via 10.1.12.2, 00:00:53, Serial1/0.12
O       5.5.5.5 [110/193] via 10.1.12.2, 00:00:53, Serial1/0.12
O       10.1.24.0 [110/128] via 10.1.12.2, 00:00:53, Serial1/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:00:53, Serial1/0.12
O       10.1.45.0 [110/192] via 10.1.12.2, 00:00:53, Serial1/0.12
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O       10.1.35.0 [110/192] via 10.1.12.2, 00:00:53, Serial1/0.12
O       10.1.34.0 [110/192] via 10.1.12.2, 00:00:53, Serial1/0.12

R1#Sh mpls int

Interface                  IP             Tunnel      BGP  Static   Operational
Serial1/0.12               No             Yes         No   No       Yes
```

On R2

```
R2#Sh ip route ospf | i O

O       1.1.1.1 [110/65] via 10.1.12.1, 00:01:33, Serial1/0.21
O       3.3.3.3 [110/65] via 10.1.23.3, 00:01:33, Serial1/0.23
O       4.4.4.4 [110/65] via 10.1.24.4, 00:01:33, Serial1/0.24
O       5.5.5.5 [110/129] via 10.1.24.4, 00:01:33, Serial1/0.24
O       10.1.45.0 [110/128] via 10.1.24.4, 00:01:33, Serial1/0.24
O       10.1.35.0 [110/128] via 10.1.23.3, 00:01:33, Serial1/0.23
O       10.1.34.0 [110/128] via 10.1.24.4, 00:01:33, Serial1/0.24

R2#Sh mpls int

Interface                  IP             Tunnel      BGP  Static   Operational
Serial1/0.21               No             Yes         No   No       Yes
Serial1/0.23               No             Yes         No   No       Yes
Serial1/0.24               No             Yes         N    No       Yes
```

On R3

```
R3#Sh ip route ospf | i O

O       1.1.1.1 [110/129] via 10.1.23.2, 00:01:56, Serial1/0.32
O       2.2.2.2 [110/65] via 10.1.23.2, 00:01:56, Serial1/0.32
O       4.4.4.4 [110/65] via 10.1.34.4, 00:01:56, Serial1/0.34
O       5.5.5.5 [110/65] via 10.1.35.5, 00:01:56, Serial1/0.35
O       10.1.12.0 [110/128] via 10.1.23.2, 00:01:56, Serial1/0.32
O       10.1.24.0 [110/128] via 10.1.34.4, 00:01:56, Serial1/0.34
O       10.1.45.0 [110/128] via 10.1.35.5, 00:01:56, Serial1/0.35

R3#Sh mpls int

Interface                  IP             Tunnel      BGP  Static   Operational
Serial1/0.32               No             Yes         No   No       Yes
Serial1/0.34               No             Yes         No   No       Yes
Serial1/0.35               No             Yes         No   No       Yes
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
R4#Sh ip route ospf | i O

O       1.1.1.1 [110/129] via 10.1.24.2, 00:02:25, Serial1/0.42
O       2.2.2.2 [110/65] via 10.1.24.2, 00:02:25, Serial1/0.42
O       3.3.3.3 [110/65] via 10.1.34.3, 00:02:25, Serial1/0.43
O       5.5.5.5 [110/65] via 10.1.45.5, 00:02:25, Serial1/0.45
O       10.1.12.0 [110/128] via 10.1.24.2, 00:02:25, Serial1/0.42
O       10.1.23.0 [110/128] via 10.1.34.3, 00:02:25, Serial1/0.43
O       10.1.35.0 [110/128] via 10.1.45.5, 00:02:25, Serial1/0.45

R4#Sh mpls int

Interface                  IP             Tunnel      BGP  Static   Operational
Serial1/0.42               No             Yes         No   No       Yes
Serial1/0.43               No             Yes         No   No       Yes
Serial1/0.45               No             Yes         NO   No       Yes
```

On R5

```
R5#Sh ip route ospf | i O

O       1.1.1.1 [110/193] via 10.1.45.4, 00:02:59, Serial1/0.54
O       2.2.2.2 [110/129] via 10.1.45.4, 00:02:59, Serial1/0.54
O       3.3.3.3 [110/65] via 10.1.35.3, 00:02:59, Serial1/0.53
O       4.4.4.4 [110/65] via 10.1.45.4, 00:02:59, Serial1/0.54
O       10.1.12.0 [110/192] via 10.1.45.4, 00:02:59, Serial1/0.54
O       10.1.24.0 [110/128] via 10.1.45.4, 00:02:59, Serial1/0.54
O       10.1.23.0 [110/128] via 10.1.35.3, 00:02:59, Serial1/0.53
O       10.1.34.0 [110/128] via 10.1.45.4, 00:02:59, Serial1/0.54

R2#Sh mpls int

Interface                  IP             Tunnel      BGP  Static   Operational
Serial1/0.53               No             Yes         No   No       Yes
Serial1/0.54               No             Yes         No   No       Yes
```

Task 2

Create a new static tunnel from R1 to R5 (5.5.5.5). Tunnel “0” should have the attributes listed in the table below and should use autoroute announce. Create a second tunnel on R2 with the same characteristics as Tunnel “0” that protects the link between R2 and R3.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The static path for the newly formed tunnel on R2 should be through the S1/0.23 next hop.

Tunnel	Bandwidth	Priority (Setup/Hold)
Tunnel 0	30	7/7

Configuring the tunnels:

On R1

```
R1(config)#interface Tunnel0
R1(config-if)#ip unnumbered Loopback0
R1(config-if)#tunnel destination 5.5.5.5
R1(config-if)#tunnel mode mpls traffic-eng
R1(config-if)#tunnel mpls traffic-eng autoroute announce
R1(config-if)#tunnel mpls traffic-eng priority 7 7
R1(config-if)#tunnel mpls traffic-eng bandwidth  30
R1(config-if)#tunnel mpls traffic-eng path-option 1 explicit name TST
R1(config-if)#tunnel mpls traffic-eng fast-reroute
R1(config-if)#exit
```

Fast Reroute enables protection of the tunnel from the Head End side. The default features of Fast Reroute are "bandwidth" and "node" protection and will be described later in the lab.

```
R1(config)#ip explicit path name TST enable
R1(cfg-ip-expl-path)#next-address 10.1.12.2
R1(cfg-ip-expl-path)#next-address 10.1.23.3
R1(cfg-ip-expl-path)#next-address 10.1.35.5
```

The tunnel should come up and use the path through R2.

```
R1#Trace 5.5.5.5

Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 10.1.12.2 [MPLS: Label 200 Exp 0] 72 msec 40 msec 36 msec
  2 10.1.23.3 [MPLS: Label 300 Exp 0] 20 msec 20 msec 24 msec
  3 10.1.35.5 20 msec *  28 msec
```

The packet is labeled across the correct interface.

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#interface Tunnel1
R2(config-if)#ip unnumbered Loopback0
R2(config-if)#tunnel destination 3.3.3.3
R2(config-if)#tunnel mode mpls traffic-eng
R2(config-if)#tunnel mpls traffic-eng autoroute announce
R2(config-if)#tunnel mpls traffic-eng bandwidth 30
R2(config-if)#tunnel mpls traffic-eng priority 7 7
R2(config-if)#tunnel mpls traffic-eng path-option 1 explicit name TST-BAK
R2(config-if)#exit

R2(config)#ip explicit-path name TST-BAK enable
R2(cfg-ip-expl-path)#next-address 10.1.24.4
R2(cfg-ip-expl-path)#next-address 10.1.34.3
```

The tunnel destination should always be the next hop router on the other side of the protected link when performing "Link Protection". The exact static path has been typed in for clarity. However, the use of the "exclude" command to prevent certain next hops from being used could have been a valid option.

```
R2(cfg-ip-expl-path)#interface s1/0.23
R2(config-subif)#mpls traffic-eng backup-path Tunnel 1
```

The interface that is being protected must be configured to allow the backup tunnel to be used. This form of a tunnel is known as "Link Protection".

Verify the Link Protected Fast Reroute Tunnel:

On R2

```
R2#Sh mpls traffic-eng fast-reroute database
```

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

LSP midpoint frr information:
LSP identifier In-label Out intf/label FRR intf/label Status
1.1.1.1 0 [4] 200 Se1/0.23:300 **Tu1**:300 **ready**

The fastreroute tunnel has already signalled the fast-reroute information to this node. It now understands what actions to take in case of a failure by leveraging RSVP signalling as a reroute mechanism.

We simulate a failure on the link between R2 and R3!!

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2(config)#int s1/0.23
R2(config-subif)#shut

R2#Sh mpls traffic-eng fast-reroute database
```

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

LSP midpoint frr information:
LSP identifier In-label Out intf/label FRR intf/label Status
1.1.1.1 0 [4] 200 Se1/0.23:300 Tu1:300 **active**

The route goes active and the traffic has a slight disruption of service.

On R1

```
R1#Trace 5.5.5.5

Type escape sequence to abort.
Tracing the route to 5.5.5.5

  1 10.1.12.2 [MPLS: Label 200 Exp 0] 64 msec 32 msec 28 msec
  2 10.1.24.4 [MPLS: Labels 400/300 Exp 0] 24 msec 24 msec 20 msec
  3 10.1.34.3 [MPLS: Label 300 Exp 0] 28 msec 16 msec 36 msec
  4 10.1.35.5 28 msec *  36 msec
```

The trace is still operational from R1. However, the path has clearly changed and is taking the backup route. Label "400" has also been added as the outer tag in addition to label "300", which represents the original tag that will be used by R3.

```
R1#Sh mpls traffic-eng tunnels | sec History:

 History:
  Tunnel:
    Time since created: 8 minutes, 12 seconds
    Time since path change: 8 minutes, 11 seconds
    Number of LSP IDs (Tun_Instances) used: 1
  Current LSP:
    Uptime: 1 minutes, 12 seconds
    Last Error: PCALC:: Explicit path has unknown address, 10.1.23.3
```

We also verify that the last LSP was indeed torn down due to a Path CALCulation error when the interface was shutdown.

Restore the broken Link on R2:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
R2(config)#int s1/0.23
R2(config-subif)#NO shut
```

Task 3

Create a tunnel named Tunnel "2" on R2. The attributes should be similar to that of tunnel on R1. The tunnel should use a fast reroute tunnel that bypasses the failure no matter what point in the path the failure takes place.

Configuring the tunnel:

On R2

```
R2(config)#interface Tunnel2
R2(config-if)#ip unnumbered Loopback0
R2(config-if)#tunnel destination 5.5.5.5
R2(config-if)#tunnel mode mpls traffic-eng
R2(config-if)#tunnel mpls traffic-eng autoroute announce
R2(config-if)#tunnel mpls traffic-eng bandwidth 30
R2(config-if)#tunnel mpls traffic-eng path-option 1 explicit name TST-BAK-2

R2(config-if)#ip explicit-path name TST-BAK-2
R2(cfg-ip-expl-path)#next-address 10.1.24.4
R2(cfg-ip-expl-path)#next-address 10.1.45.5

R2(cfg-ip-expl-path)#int s1/0.23
R2(config-subif)#mpls traffic-eng backup-path Tunnel 2
```

Here the tunnel destination represents the loopback interface of R5 instead of R3. In this case, "Node-Protection" requires that the termination be at least one hop beyond the next hop router, refferred to as NNHOP. The interface is again protected but by a different tunnel.

Verifying Tunnel operation:

```
R2#Sh mpls traffic-eng fast-reroute database
```

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

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

LSP midpoint frr information:
LSP identifier In-label Out intf/label FRR intf/label Status
1.1.1.1 0 [26] 201 Se1/0.23:301 **Tu2**:implicit-nul **ready**

The tunnel is ready but the tunnel it is taking has changed to be tunnel "2". The router will always see a path that protects a NODE as being superior. This type of protection would protect anywhere in the path, not just a link failure.

```
R2(config)#int s1/0.23
R2(config-subif)#shut
```

Simulating the error again.

```
R2#Sh mpls traffic-eng fast-reroute database
```

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

LSP midpoint frr information:
LSP identifier In-label Out intf/label FRR intf/label Status
1.1.1.1 0 [26] 201 Se1/0.23:301 Tu2:implicit-nul active

The route is now active to reroute traffic onto the different path specified.

```
R1#Trace 5.5.5.5

Type escape sequence to abort.
Tracing the route to 5.5.5.5

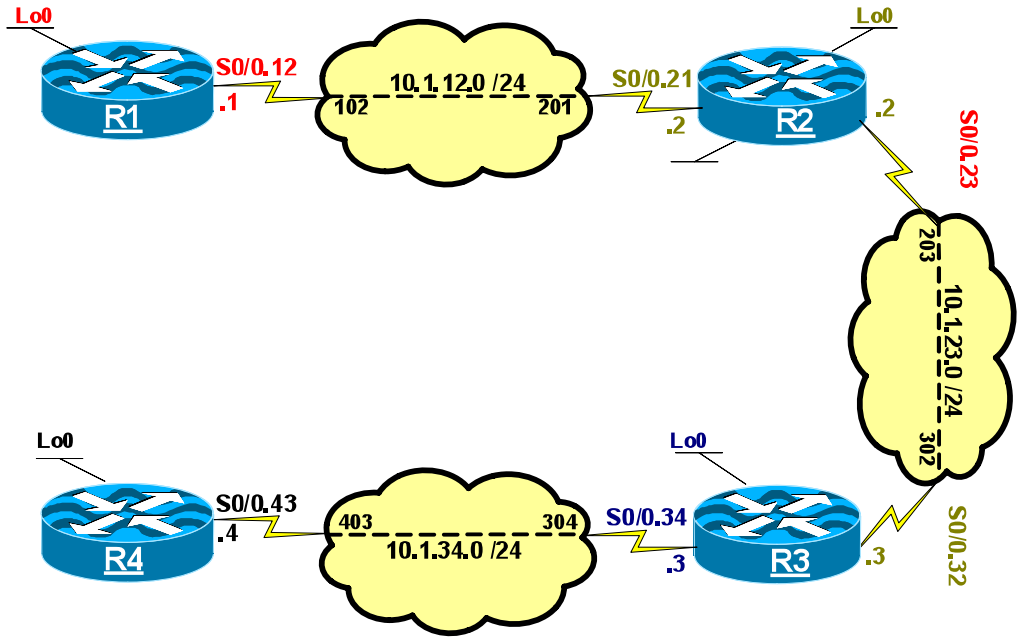
  1 10.1.12.2 [MPLS: Label 201 Exp 0] 412 msec 32 msec 20 msec
  2 10.1.24.4 [MPLS: Label 401 Exp 0] 32 msec 20 msec 16 msec
  3 10.1.45.5 28 msec *  24 msec
```

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 5 –MPLS QOS



IP Addressing:

Router	Interface	IP address
R1	Lo0 S0/0.12	1.1.1.1 /32 10.1.12.1 /24
R2	Lo0 S0/0.21 S0/0.23	2.2.2.2 /32 10.1.12.2 /24 10.1.23.2/ 24
R3	Lo0 S0/0.32 S0/0.34	3.3.3.3/ 32 10.1.23.3 /24 10.1.34.3. 24
R4	Lo0 S0/0.43	4.4.4.4 /32 10.1.34.4 /24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab Setup:

The lab has all frame relay PVC's configured in a point-to-point manner.

Task 1

Configure an OSPF with a process of "1" on all routers. The R2, R3 and R4 routers must be prepared to support MPLS-TE. All RSVP interfaces should be configured to allow a global bandwidth utilization of 100K.

On R1

```
R1(config)#router ospf 1
R1(config-router)#network 0.0.0.0 0.0.0.0 area 0
```

On R2

```
R2(config)#mpls traffic-eng tunnels
R2(config)#mpls label range 200 299

R2(config)#router ospf 1
R2(config-router)#network 0.0.0.0 0.0.0.0 area 0
R2(config-router)#mpls traffic-eng area 0
R2(config-router)#mpls traffic-eng router-id loopback 0

R2(config)#interface S0/0.23
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth 100
```

On R3

```
R3(config)#mpls traffic-eng tunnels
R3(config)#mpls label range 300 399

R3(config)#router ospf 1
R3(config-router)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#mpls traffic-eng area 0
R3(config-router)#mpls traffic-eng router-id loopback 0

R3(config)#interface S0/0.32
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth 100

R3(config-if)#interface S0/0.34
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth 100
```

On R4

```
R4(config)#mpls traffic-eng tunnels
R4(config)#mpls label range 400 499

R4(config)#router ospf 1
R4(config-router)#network 0.0.0.0 0.0.0.0 area 0
R4(config-router)#mpls traffic-eng area 0
R4(config-router)#mpls traffic-eng router-id loopback 0

R4(config-if)#interface S0/0.43
R4(config-if)#mpls traffic-eng tunnels
R4(config-if)#ip rsvp bandwidth 100
```

Verify the Configuration

On R1

```
R1#Sh ip route ospf | i O

O       2.2.2.2 [110/65] via 10.1.12.2, 00:00:39, Serial0/0.12
O       3.3.3.3 [110/129] via 10.1.12.2, 00:00:39, Serial0/0.12
O       4.4.4.4 [110/193] via 10.1.12.2, 00:00:39, Serial0/0.12
O       10.1.23.0 [110/128] via 10.1.12.2, 00:00:39, Serial0/0.12
O       10.1.34.0 [110/192] via 10.1.12.2, 00:00:39, Serial0/0.12
```

On R2

```
R2#Sh ip route ospf | i O

O       1.1.1.1 [110/65] via 10.1.12.1, 00:00:50, Serial0/0.21
O       3.3.3.3 [110/65] via 10.1.23.3, 00:00:50, Serial0/0.23
O       4.4.4.4 [110/129] via 10.1.23.3, 00:00:50, Serial0/0.23
O       10.1.34.0 [110/128] via 10.1.23.3, 00:00:50, Serial0/0.23

R2#Sh mpls int

Interface                  IP             Tunnel      Operational
Serial0/0.23               No             Yes         Yes
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh ip route ospf | i O

O       1.1.1.1 [110/129] via 10.1.23.2, 00:01:11, Serial0/0.32
O       2.2.2.2 [110/65] via 10.1.23.2, 00:01:11, Serial0/0.32
O       4.4.4.4 [110/65] via 10.1.34.4, 00:01:11, Serial0/0.34
O       10.1.12.0 [110/128] via 10.1.23.2, 00:01:11, Serial0/0.32

R3#Sh mpls int

Interface                  IP             Tunnel      Operational
Serial0/0.32               No             Yes         Yes
Serial0/0.34               No             Yes         Yes
```

On R4

```
R4#Sh ip route ospf | i O

O       1.1.1.1 [110/193] via 10.1.34.3, 00:02:14, Serial0/0.43
O       2.2.2.2 [110/129] via 10.1.34.3, 00:02:14, Serial0/0.43
O       3.3.3.3 [110/65] via 10.1.34.3, 00:02:14, Serial0/0.43
O       10.1.12.0 [110/192] via 10.1.34.3, 00:02:14, Serial0/0.43
O       10.1.23.0 [110/128] via 10.1.34.3, 00:02:14, Serial0/0.43

R4#Sh mpls int

Interface                  IP             Tunnel      Operational
Serial0/0.43               No             Yes         Yes
```

Task 2

Create a tunnel called “Tunnel”0” from R2 to R4. This tunnel must have the following characteristics: setup and hold of “1”, bandwidth of 40K, dynamic setup.

On R2

```
R2(config)#int tunnel 0
R2(config-if)#ip unnumbered lo0
R2(config-if)#tunnel destination 4.4.4.4
R2(config-if)#tunnel mode mpls traffic-eng
R2(config-if)#tunnel mpls traffic-eng priority 1 1
R2(config-if)#tunnel mpls traffic-eng bandwidth 40
R2(config-if)#tunnel mpls traffic-eng path-option 1 dynamic
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Verify the tunnel:

On R2

```
R2#Sh ip rsvp reservation

To          From        Pro  DPort  Sport  Next Hop     I/F        Fi  Serv    BPS
4.4.4.4     2.2.2.2     0    0      1      10.1.23.3    Se0/0.23   SE  LOAD    40K

R2#Sh mpls traffic-eng link-management admission-control

System Information::
  Tunnels Count:       1
  GMPLS Tunnels Count: 0
  Tunnels Selected:    1
TUNNEL ID        UP IF    DOWN IF   PRIORITY  STATE           BW (kbps)
2.2.2.2 0_1      -        Se0/0.23  1/1       Resv Admitted   40          RG

R2#Sh mpls traffic-eng link-management advertisements | b Res

Res. Global BW:     100 kbits/sec
  Res. Sub BW:      0 kbits/sec
  Downstream::
                                    Global Pool      Sub Pool
                                    -------------    ----------
    Reservable Bandwidth[0]:        100              0 kbits/sec
    Reservable Bandwidth[1]:        60               0 kbits/sec
    Reservable Bandwidth[2]:        60               0 kbits/sec
    Reservable Bandwidth[3]:        60               0 kbits/sec
    Reservable Bandwidth[4]:        60               0 kbits/sec
    Reservable Bandwidth[5]:        60               0 kbits/sec
    Reservable Bandwidth[6]:        60               0 kbits/sec
    Reservable Bandwidth[7]:        60               0 kbits/sec
  Attribute Flags:      0x00000000
```

The global pool has been allowed up to 100 kbps. Due to the tunnel priority, all other priorites have also adjusted the available bandwidth to be 60.

Task 3

Create a tunnel called “Tunnel”1” that has similar characteristics to that of tunnel “0”.
The only required change in the configuration will be to allow for bandwidth guarantees

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

to be possible for certain tunnel priorities above “3”. The guarantee for Tunnel “1” should be 25K.”

Configuring DS-TE

On R2

```
R2(config)#int tunnel 1
R2(config-if)#ip unnumbered Loopback0
R2(config-if)#tunnel destination 4.4.4.4
R2(config-if)#tunnel mode mpls traffic-eng
R2(config-if)#tunnel mpls traffic-eng autoroute announce
R2(config-if)#tunnel mpls traffic-eng priority 4 4
R2(config-if)#tunnel mpls traffic-eng bandwidth sub-pool 25
R2(config-if)#tunnel mpls traffic-eng path-option 1 dynamic
```

The Tunnel Head-End can choose between “bandwidth” or “subpools” as an option but not both. Sub-pools allow for more granular usage of RSVP bandwidth reservations. Traffic that is routed down a tunnel can still experience congestion. The goal is to maintain some kind of banwidth management that fits into the current QOS model.

```
R2(config)#interface s0/0.23
R2(config-subif)#ip rsvp bandwidth 100 sub-pool 40
```

ON R3

```
R3(config)#interface s0/0.32
R3(config-subif)#ip rsvp bandwidth 100 sub-pool 40

R3(config-subif)#interface s0/0.34
R3(config-subif)#ip rsvp bandwidth 100 sub-pool 40
```

ON R4

```
R4(config)#interface s0/0.43
R4(config-subif)#ip rsvp bandwidth 100 sub-pool 40
```

The entire path must have subpools available in RSVP for ANY tunnel that uses sub-pools as an option to operate correctly. R4 does not need the configuration as it represents the last hop. The configuration is there for completeness.

Verify the change:

On R2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R2#Sh mpls traffic-eng link-management advertisements | b Res
```

Res. Global BW: 100 kbits/sec
Res. Sub BW: 40 kbits/sec
Downstream::

	Global Pool	Sub Pool
Reservable Bandwidth[0]:	100	**40 kbits/sec**
Reservable Bandwidth[1]:	60	**40 kbits/sec**
Reservable Bandwidth[2]:	60	**40 kbits/sec**
Reservable Bandwidth[3]:	60	**40 kbits/sec**
Reservable Bandwidth[4]:	**35**	**15 kbits/sec**
Reservable Bandwidth[5]:	**35**	**15 kbits/sec**
Reservable Bandwidth[6]:	**35**	**15 kbits/sec**
Reservable Bandwidth[7]:	**35**	**15 kbits/sec**

Attribute Flags: 0x00000000

The tunnel reserves the bandwidth that has been reserved for the sub-pools as well as the global pool. The same does not hold true for the global reservations. They can only reserve bandwidth from the Global Pool only. The output above also shows how the bandwidth from a certain priority will reserve from all priorities that are inferior to it.

Task 4

All pings will be generated from R1 with a TOS of 184, 96 or 32. Create a service-policy on R2 that matches on a value other than DSCP or IPP. No filters are permitted. A previous marking must be matched. Create the same policy on R3 to verify the markings.

Configure QOS to match and mark packets:

On R2

```
R2(config)#class-map TST-5
R2(config-cmap)#match mpls experimental 5
R2(config-cmap)#exit

R2(config)#class-map TST-3
R2(config-cmap)#match mpls experimental 3
R2(config-cmap)#exit
```

We can match the experimental field due to the fact that the 3 most significant bits in the TOS field are

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

automatically copied into the experimental field.

```
R2(config)#policy-map MARK-TST
R2(config-pmap)#class TST-5
R2(config-pmap-c)#priority percent 20
R2(config-pmap-c)#exit
R2(config-pmap)#class TST-3
R2(config-pmap-c)#bandwidth percent 10
R2(config-pmap-c)#exit

R2(config)#int s0/0
R2(config-subif)#service-policy output MARK-TST
```

On R3

```
R3(config)#class-map TST-5
R3(config-cmap)#match mpls experimental 5
R3(config-cmap)#exit

R3(config)#class-map TST-3
R3(config-cmap)#match mpls experimental 3
R3(config-cmap)#exit

R3(config)#policy-map MARK-TST
R3(config-pmap)#class TST-5
R3(config-pmap-c)#set mpls experimental 5
R3(config-pmap-c)#exit
R3(config-pmap)#class TST-3
R3(config-pmap-c)#set mpls experimental 3
R3(config-pmap-c)#exit

R3(config)#int s0/0.32
R3(config-subif)#service-policy input MARK-TST
```

This policy is only created to verify the marking is being performed by R2

Verify the markings:

On R1

```
R1#Ping
Protocol [ip]:
Target IP address: 4.4.4.4
Repeat count [5]: 150
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 1.1.1.1
Type of service [0]: 160
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 150, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!
Success rate is 100 percent (150/150), round-trip min/avg/max = 4/18/72 ms

R1#Ping
Protocol [ip]:
Target IP address: 4.4.4.4
Repeat count [5]: 99
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 1.1.1.1
Type of service [0]: 96
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 99, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (99/99), round-trip min/avg/max = 8/21/356 ms
```

On R2

```
R2#Sh policy-map int s0/0

 Serial1/0

 Service-policy output: MARK-TST
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    queue stats for all priority classes:
      Queueing
      queue limit 64 packets
      (queue depth/total drops/no-buffer drops) 0/0/0
      (pkts output/bytes output) 150/16200

    Class-map: TST-5 (match-all)
      150 packets, 15600 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: mpls experimental topmost 5
      Priority: 20% (308 kbps), burst bytes 7700, b/w exceed drops: 0

   Class-map: TST-3 (match-all)
      99 packets, 10296 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: mpls experimental topmost 3
      Priority: 10% (154 kbps), burst bytes 3850, b/w exceed drops: 0
```

—-text omitted—-

The packets were marked and placed in the proper queues.

On R3

```
R3#Sh policy-map int
 Serial1/0

  Service-policy input: MARK-TST

Class-map: TST-5 (match-all)
      150 packets, 16200 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: mpls experimental topmost 5
      QoS Set
        mpls experimental imposition 5
          Packets marked 150

Class-map: TST-3 (match-all)
      99 packets, 10692 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: mpls experimental topmost 3
      QoS Set
        mpls experimental imposition 3
          Packets marked 99
```

—-text omitted—-

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Modify the traffic at Experimental value of 3 to be remarked to 1 on Tunnel "1", MPLS TE bandwidth should only be flooded when traffic is measured rising at 20% and 60 %. The bandwidth info should also be flooded when traffic decreases at 40% and 80 %.

On R2

```
R2(config)#policy-map MARK-TST
R2(config-pmap)#class TST-3
R2(config-pmap-c)#set mpls experimental topmost 1
```

The mpls experimental bit cannot be set on an outbound policy with the "imposition" keyword. The "topmost" keyword must be used instead.

On R3

```
R3(config)#class-map REMARK-TST
R3(config-cmap)#match mpls exp 1

R3(config-cmap)#policy-map MARK-TST
R3(config-pmap)#class REMARK-TST
R3(config-pmap-c)#set mpls exper 1
```

Verify the remarking:

On R1

```
R1#Ping
Protocol [ip]:
Target IP address: 4.4.4.4
Repeat count [5]: 101
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface: 1.1.1.1
Type of service [0]: 96
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 101, 100-byte ICMP Echos to 4.4.4.4, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
!!!
```

4	0	0	0	40
5	0	0	0	40
6	0	0	0	40
7	0	0	0	40

Downstream Sub Pool Bandwidth Information (kbits/sec):

KEEP PRIORITY	BW HELD	BW TOTAL HELD	BW LOCKED	BW TOTAL LOCKED
0	0	0	0	0
1	0	0	0	0
2	0	0	0	0
3	0	0	0	0
4	0	0	**25**	**25**
5	0	0	0	**25**
6	0	0	0	**25**
7	0	0	0	**25**

The bandwidth is allocated to the subpool. Notice how the "MAX Res Sub BW" field has been changed to 40k which is the reserved bandwidth for the tunnel and the "BW Descriptors" field still has 1 descriptor but also shows that the descriptor includes the 1 subpool configured. This feature allows the tunnel to take advantage of a Modular QOS Policy that has LLQ enabled for priority bandwidth. The bandwidth that is provisioned in the Policy will utilize 100% of the subpool bandwidth out of the interface that also has the sub pool command towards the destination.

```
R2#Sh mpls traffic-eng link-management bandwidth-allocation | s Up | Down Thre
```

Up Thresholds: 15 30 45 60 75 80 85 90 95 96 97 98 99 100 (default)
Down Thresholds: 100 99 98 97 96 95 90 85 80 75 60 45 30 15 (default)

The thresholds show when the bandwidth information will be flooded into TE calculations for either UP % or Down %.

On R2

```
R2(config-if)#int S0/0
R2(config-if)#mpls traffic-eng flooding  thresholds down 40 80
R2(config-if)#mpls traffic-eng flooding  thresholds up 20 60
```

The default thresholds have many changes imposed at different levels. If we wish for a value to remain, the adjustment shoud contain the original value.

```
R2#Sh mpls traffic-eng link-management bandwidth-allocation | s Up | Down Thre
```

Up Thresholds: **20 60**
Down Thresholds: **80 40**

The other thresholds are gone but what we have configured has remained.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

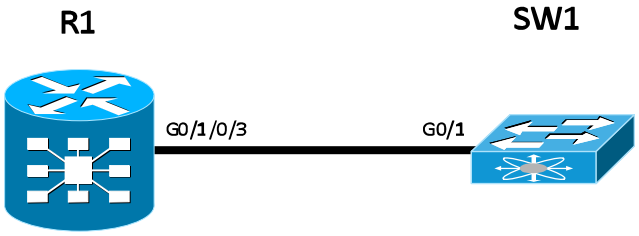
Paul Negron
CCIE #14856
SP

IOS-XR

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

LAB-1
IOS-XR Quick Reference



Lab Setup:

Router	Interface	IP address
R1	Lo0 G0/1/0/3	1.1.1.1/32 10.1.12.1 /24
SW1	Lo0 F0/0	2.2.2.2 /32 10.1.12.2 /24

Task 1

Configure a hostname on the IOS-XR platform.

On R1

```
12A con0/3/CPU0 is now available

Press RETURN to get started.

User Access Verification

Username:
```

When looging into an IOS-XR platform, The prompt does not enter "user mode" as it does with IOS. A password is required.

```
Username: cisco
Password: cisco
RP/0/3/CPU0:Mar 10 16:13:15.562 : exec[65696]: %SECURITY-login-6-
AUTHEN_SUCCESS : Successfully authenticated user 'cisco' from 'console'
on 'con0_3_CPU0'
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/3/CPU0:ios#
```

The default username is "ios" but the device will always show you which RP and slot you are logged into. In this case, it is the RP located in slot 3. The 12000 series router displays the location in te following format: "*module*/*slot*/*CPU*:name"

```
RP/0/3/CPU0:ios#Conf
RP/0/3/CPU0:ios(config)#hostname R1
```

Notice how the router did not apply the configuration right away.

```
RP/0/3/CPU0:ios(config)#Show configuration

Building configuration...
!! IOS XR Configuration 0.0.0
hostname R1
end
```

Here we see the configuration is stored in a parser that has only made sure that the sytax check has passed but the symantex check. This location is called the "target" configuration. We can make changes to the configuration or clear its contents without affecting the running configuration. This is also called the First stage of configuration. Also, notice how the "show configuration" command was executed in global configuration mode.

```
RP/0/3/CPU0:ios(config)#Commit
RP/0/3/CPU0:Mar 10 16:20:27.844 : config[65731]: %MGBL-CONFIG-6-DB_COMMIT :
Configuration committed by user 'cisco'. Use 'show configuration commit changes
1000000020' to view the changes.
RP/0/3/CPU0:R1(config)#
```

In order for the configuration to be a part of the running configuration, a "commit" must be issued. When a commit is issued, several actions take place. 1) A copy of the config is stored as a binary on the disk. 2) The target = running configuration. 3) The config is stored in a commit list to be viewed or to allow a rollback action just in case it caused any probems. The filed has a default label of a long indexed number, in this case, "1000000020".

```
RP/0/3/CPU0:R1(config)#Sh config commit list

SNo. Label/ID    User     Line       Client     Time Stamp
~~~~ ~~~~~~~~~~  ~~~~     ~~~~       ~~~~~~~    ~~~~~~~~~~~~
1    1000000020  cisco    con0_3_CPU  CLI       Sat Mar 10 16:20:27 2012
2    1000000019  cisco    con0_3_CPU  CLI       Sat Mar 10 16:19:01 2012
3    1000000018  cisco    con0_3_CPU  CLI       Thu Mar  8 07:00:27 2012
*** text omitted ***
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

When we commited the configuration it caused the "1000000020" change to be created. Since it was just recently commited, it appears at the top of the list ahead of the previous change.

```
RP/0/3/CPU0:12A(config)#exit

RP/0/3/CPU0:R1#Sh config commit changes 1000000020

Building configuration...
!! IOS XR Configuration 0.0.0
hostname 12A
end
```

Because we are viewing the changes with the "comitt" option, the individual commands that were stored in the target config just before you commited them are seen.

```
RP/0/3/CPU0:R1#Sh configuration rollback changes 1000000020
Sat Mar 10 22:25:09.652 UTC
Building configuration...
!! IOS XR Configuration 0.0.0
hostname ios
end
```

By using the "rollback" option, the output shows what would happen if you rolled back that change. In most cases, the "no" keyword would be inserted to negate the change. Sometimes a default may show up as the action.

```
RP/0/3/CPU0:R1(config)#exit

RP/0/3/CPU0:12A#rollback  configuration last 1

Loading Rollback Changes.
Loaded Rollback Changes in 1 sec
Committing.
1 items committed in 1 sec (0)items/sec
Updating.RP/0/3/CPU0:Mar 10 16:23:03.572 : config_rollback[65731]: %MGBL-CONFIG-6-
DB_COMMIT : Configuration committed by user 'cisco'. Use 'show configuration commit
changes 1000000021' to view the changes.
```

Rolling back a configuration to the last change will cause another rollback point to be created with a unique label.

```
Updated Commit database in 1 sec
Configuration successfully rolled back 1 commits.

RP/0/3/CPU0:ios#
RP/0/3/CPU0:ios(config)#load commit changes 1000000020
Building configuration...
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
Loading.
48 bytes parsed in 1 sec (47)bytes/sec
```

We can commit the individual change if needed as well.

```
RP/0/3/CPU0:ios(config)#Sh config

Building configuration...
!! IOS XR Configuration 0.0.0
hostname R1
end
```

When loading a configuration, it is stored in the target but still needs to be commited. It is almost like copying from a notepad and pasting to the console. The "clear" command in this mode would clear the target configuration as if you did not type anything.

```
RP/0/3/CPU0:ios(config)#commit
RP/0/3/CPU0:Mar 10 16:25:49.783 : config[65731]: %MGBL-CONFIG-6-DB_COMMIT :
Configuration committed by user 'cisco'. Use 'show configuration commit changes
1000000022' to view the changes.
RP/0/3/CPU0:R1(config)#
```

After commiting again, another unique label is created. A lot of time can be spent just playing with the "commit" and "rollback" feature but it is important in a real production environment to make safe changes.

Task 2

Configure IP connectivity between the 3400 and the 12000 Series Router using the table presented in the lab setup.

On SW1

```
SW1(config)#int lo0
SW1(config-if)#ip add 2.2.2.2 255.255.255.255

SW1(config-if)#int g0/1
SW1(config-if)#No switchport
SW1(config-if)#ip address 10.1.12.2 255.255.255.0
SW1(config-if)#No shut
```

The standard IOS way of doing things.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
RP/0/3/CPU0:R1(config)#int lo0
RP/0/3/CPU0:R1(config-if)#ipv4 address 1.1.1.1/32

RP/0/3/CPU0:R1(config-if)#int g0/1/0/3
RP/0/3/CPU0:R1(config-if)#ipv4 address 10.1.12.1/24
RP/0/3/CPU0:R1(config-if)#No shut
RP/0/3/CPU0:R1(config-if)#end
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:yes
```

Here we see a few changes woth IOS-XR. The protocol must always be specified (ipv4). The mask can be configured with a decimal or bit mask format. Also, commands can often be built in the same line. AN example of thi is when applying an access-list. "int g0/1/0/3 ipv4 address 10.1.12.2/24" is acceptable but adding "no shut" on the same line would not.

Whe tried to end the configuration like in IOS but we did not commit first.
Yes = commit the changes and continue back to "EXEC" mode.
No = do not commit the changes and continue back to "EXEC" mode
Cancel = "I changed my mind so do nothing. You also stay in that mode.

Verify Connectivity:

On SW1

```
SW1#Sh run int g0/1
Building configuration...

Current configuration : 102 bytes
!
interface GigabitEthernet0/1
 port-type nni    "This command is covered in the switching section"
 no switchport
 ip address 10.1.12.2 255.255.255.0
end
```

On R1

```
RP/0/3/CPU0:R1#Sh run int g0/1/0/3

interface GigabitEthernet0/1/0/3
 ipv4 address 10.1.12.1 255.255.255.0

RP/0/3/CPU0:R1#Ping 10.1.12.2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.12.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 2/2/3 ms
```

Task 3

Configure IP connectivity between the 3400 and the 12000 Series Router using the table presented in the lab setup.

On SW1

```
SW1(config)#ip routing
SW1(config)#ip route 1.0.0.0 255.0.0.0 10.1.12.1
```

On R1

```
RP/0/3/CPU0:R1(config)#router static
RP/0/3/CPU0:R1(config-static)#address ipv4 unicast
RP/0/3/CPU0:R1(config-static-afi)#2.0.0.0/8 10.1.12.2
```

This is very different than IOS. The IOS-XR code uses protocol based CLI. If the protocol is "address-family aware, you have an option to use it. Static routes require the address-family option. In this case, it is for ipv4.

Verify Connectivity:

On SW1

```
SW1#Sh run | i ip route

 no ip route-cache cef
 no ip route-cache
ip route 1.0.0.0 255.0.0.0 10.1.12.1
```

The "|" option can be very useful to parse portions of configuration and outout. Sometimes more filtering must be performed to take out unwanted text.

```
SW1#Sh ip route static
S    1.0.0.0/8 [1/0] via 10.1.12.1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
RP/0/3/CPU0:R1#Sh run router static

router static
 address-family ipv4 unicast
  2.0.0.0/8 10.1.12.2
```

The "show run" command ha smore options for fitering than IOS. We can get away without using the "|" command unless necessary.

```
RP/0/3/CPU0:R1#Sh route static

S    2.0.0.0/8 [1/0] via 10.1.12.2, 00:08:11
```

The "ip" keyword has been removed as an option for verifying routing information. The default is "ipv4". We could have typed "show route ipv4". This is fairly consistent in XR code. Other examples would be "show bgp" or "show ospf"

```
RP/0/3/CPU0:R1#Ping 2.2.2.2 source lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 22/25/34 ms
```

Task 4

Remove the static routig and configure OSPF as the dynamic routing protocol.

On SW1

```
SW1(config)#No ip route 1.0.0.0 255.0.0.0 10.1.12.1

SW1(config)#router ospf 1
SW1(config-router)#netw 10.1.12.0 0.0.0.255 area 0

SW1(config-router)#int lo0
SW1(config-if)#ip ospf 1 area 0

SW1(config)#router ospf 1
SW1(config-router)#passive-interface lo0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Here we see the classic process based and interface based OSPF configurations. The "interface" based command eliminated the need for "network" comand but was still a challenge to verify the OSPF configuration easily.

On R1

```
RP/0/3/CPU0:R1(config)#router ospf TST
RP/0/3/CPU0:R1(config-ospf)#router-id 0.0.0.1
RP/0/3/CPU0:R1(config-ospf)#area 0
RP/0/3/CPU0:R1(config-ospf)#mtu enable
RP/0/3/CPU0:R1(config-ospf-ar)#interface g0/1/0/3
RP/0/3/CPU0:R1(config-ospf-ar-if)#exit

RP/0/3/CPU0:R1(config-ospf-ar)#interface lo0
RP/0/3/CPU0:R1(config-ospf-ar-if)#passive
```

The IOS-XR configuration is quite scalable and allows for inheritance of commands that is more consistence. For example, hello timers and authentication can be performed at the process, area or interface level withing the OSPF context. If cost is configured under the area level, ALL interfaces inherit that cost. If a cost is specified under the interface within the ospf configuration, it overrides the cost at all levels above it. Since the iterface on a 3400 NNI port defaults to 9000 as the MTU and the 12000 defaulted to using an MTU if 1514, the mtu is set to enable under the area which will igmore the MTU mismatch under ospf to allow the peering to go to the full state. All of the interface inherit this and can be overridden at the interface level.

Verifying the Configuration:

SW1

```
SW1#Sh ip route ospf | i O

O       1.1.1.1 [110/2] via 10.1.12.1, 00:05:36, GigabitEthernet0/1
```

On R1

```
RP/0/3/CPU0:R1#Sh run router ospf

router ospf TST
 router-id 0.0.0.1
 area 0
  mtu-ignore enable
  interface Loopback0
   passive enable
  !
 interface GigabitEthernet0/1/0/3
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/3/CPU0:R1#Sh ospf neighbor

* Indicates MADJ interface

Neighbors for OSPF TST

Neighbor ID   Pri   State      Dead Time   Address      Interface
0.0.0.2       1     FULL/DR    00:00:38    10.1.12.2    GigabitEthernet0/1/0/3
    Neighbor is up for 00:04:06

Total neighbor count: 1

RP/0/3/CPU0:R1#Sh route ospf

2.2.2.2/32 [110/2] via 10.1.12.2, 00:06:14, GigabitEthernet0/1/0/3

RP/0/3/CPU0:R1#Ping 2.2.2.2 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/2/4 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

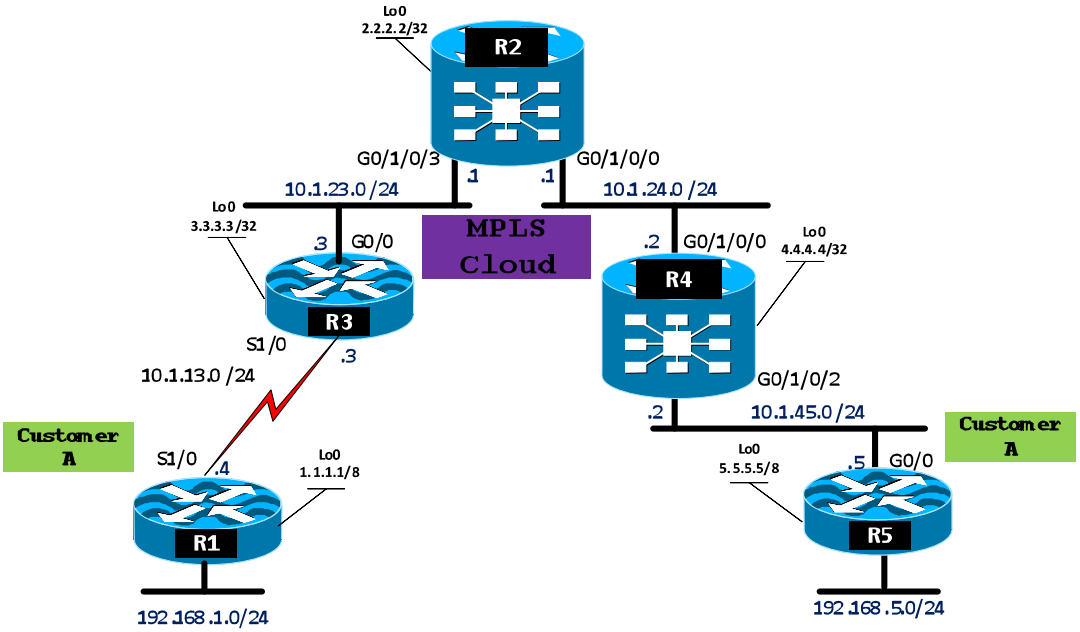
Paul Negron
CCIE #14856
SP

IOS-XR & Layer 3 VPNs

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 2
Static & RIPv2 routing in a VPN



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100 and R5 in VLAN 500.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the rest of the routers according to the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

On R2

```
RP/0/5/CPU0:R2(config)#router ospf TST
RP/0/5/CPU0:R2(config-ospf)#area 0
RP/0/5/CPU0:R2(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R2(config-ospf-ar-if)#passive
RP/0/5/CPU0:R2(config-ospf-ar-if)#commit
RP/0/5/CPU0:R2(config-ospf-ar-if)#end
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
RP/0/5/CPU0:R2(config)#router ospf TST
RP/0/5/CPU0:R2(config-ospf)#area 0
RP/0/5/CPU0:R2(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R2(config-ospf-ar-if)#passive
RP/0/5/CPU0:R2(config-ospf-ar-if)#commit
RP/0/5/CPU0:R2(config-ospf-ar-if)#end
```

To verify the configuration:

On R3

```
R3#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O        2.2.2.2 [110/2] via 10.1.23.2, 00:07:28, GigabitEthernet0/0
O        4.4.4.4 [110/3] via 10.1.23.2, 00:01:59, GigabitEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O        10.1.24.0/24 [110/2] via 10.1.23.2, 00:07:28, GigabitEthernet0/0
```

On R2

```
RP/0/5/CPU0:R2#Sh ip route ospf | i O

O    3.3.3.3/32 [110/2] via 10.1.23.3, 00:05:44, GigabitEthernet0/1/0/3
O    4.4.4.4/32 [110/2] via 10.1.24.4, 00:00:18, GigabitEthernet0/1/0/0
```

On R4

```
RP/0/5/CPU0:R4#Sh ip route ospf | i O

O    2.2.2.2/32 [110/2] via 10.1.24.2, 00:35:03, GigabitEthernet0/1/0/0
O    3.3.3.3/32 [110/3] via 10.1.24.2, 00:35:03, GigabitEthernet0/1/0/0
O    10.1.23.0/24 [110/2] via 10.1.24.2, 00:35:03, GigabitEthernet0/1/0/0
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 65200 – 65299
R3 – 65300 – 65399
R4 – 65400 – 65499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 65300 65399

R3(config)#int g0/0
R3(config-if)#MPLS IP
```

On R2

```
RP/0/5/CPU0:R2(config)#mpls label range 65200 65299
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R2(config)#mpls ldp
RP/0/5/CPU0:R2(config-ldp)#router-id 2.2.2.2
RP/0/5/CPU0:R2(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ldp-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ldp-if)#commit
RP/0/5/CPU0:R2(config-ldp-if)#end
```

On R4

```
RP/0/5/CPU0:R4(config)#mpls label range 65400 65499

RP/0/5/CPU0:R4(config)#mpls ldp
RP/0/5/CPU0:R4(config-ldp)#router-id 4.4.4.4
RP/0/5/CPU0:R4(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ldp-if)#commit
RP/0/5/CPU0:R4(config-ldp-if)#end
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Show mpls interface

Interface                  LDP      Tunnel   Enabled
-------------------------- -------- -------- --------
GigabitEthernet0/1/0/0     Yes      No       Yes

RP/0/5/CPU0:R4#Show mpls ldp neighbor

Peer LDP Identifier: 2.2.2.2:0
  TCP connection: 2.2.2.2:646 - 4.4.4.4:14775
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 11/9; Downstream-Unsolicited
  Up time: 00:01:34
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    2.2.2.2          10.1.23.2        10.1.24.2

RP/0/5/CPU0:R4#Sh mpls ldp discovery detail

Local LDP Identifier: 4.4.4.4:0
Discovery Sources:
  Interfaces:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    GigabitEthernet0/1/0/0 (0x2000300) : xmit/recv
      Source address: 10.1.24.4; Transport address: 4.4.4.4
      Hello interval: 5 sec (due in 917 msec)
      Quick-start: Enabled
      LDP Id: 2.2.2.2:0
          Source address: 10.1.24.2; Transport address: 2.2.2.2
          Hold time: 15 sec (local:15 sec, peer:15 sec)
                     (expiring in 12 sec)

RP/0/5/CPU0:R4#Show mpls label range

Range for dynamic labels: Min/Max: 65400/65499
```

The default range for the labels is 16000 to 1048575 on this platform. These number ranges are software and platform specific. Over 353,000 labels are supported on 6500's as of the writing of this work book.

On R2

```
RP/0/5/CPU0:R2#Sh mpls interface

Interface                  LDP      Tunnel   Enabled
-------------------------- -------- -------- --------
GigabitEthernet0/1/0/0     Yes      No       Yes
GigabitEthernet0/1/0/3     Yes      No       Yes

RP/0/5/CPU0:R2#Sh mpls ldp neighbor

Peer LDP Identifier: 3.3.3.3:0
  TCP connection: 3.3.3.3:60068 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 17/18; Downstream-Unsolicited
  Up time: 00:08:14
  LDP Discovery Sources:
    GigabitEthernet0/1/0/3
  Addresses bound to this peer:
    3.3.3.3          10.1.13.3        10.1.23.3

Peer LDP Identifier: 4.4.4.4:0
  TCP connection: 4.4.4.4:14775 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 15/17; Downstream-Unsolicited
  Up time: 00:06:42
  LDP Discovery Sources:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    4.4.4.4          10.1.24.4        10.1.45.4        10.1.46.4

RP/0/5/CPU0:R2#Sh mpls ldp discovery

Local LDP Identifier: 2.2.2.2:0
Discovery Sources:
  Interfaces:
    GigabitEthernet0/1/0/0 : xmit/recv
      LDP Id: 4.4.4.4:0, Transport address: 4.4.4.4
          Hold time: 15 sec (local:15 sec, peer:15 sec)

    GigabitEthernet0/1/0/3 : xmit/recv
      LDP Id: 3.3.3.3:0, Transport address: 3.3.3.3
          Hold time: 15 sec (local:15 sec, peer:15 sec)

RP/0/5/CPU0:R2#Sh mpls label range

Range for dynamic labels: Min/Max: 65200/65299
```

On R3

```
R3#Sh mpls interfaces

Interface              IP            Tunnel   BGP Static Operational
GigabitEthernet0/0     Yes (ldp)     No       No  No     Yes

R3#Show mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.23184
        State: Oper; Msgs sent/rcvd: 13/12; Downstream
        Up time: 00:04:01
        LDP discovery sources:
          GigabitEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          10.1.23.2       10.1.24.2       2.2.2.2

R3#Show mpls ldp discovery all

Local LDP Identifier:
    3.3.3.3:0
    Discovery Sources:
    Interfaces:
        GigabitEthernet0/0 (ldp): xmit/recv
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
            LDP Id: 2.2.2.2:0

R3#Show mpls label range

Downstream Generic label region: Min/Max label: 65300/65399
```

To verify the LFIB of these routers:

NOTE:

The labels displayed in the output of this lab may differ from your lab but the principle remains the same.

On R3

```
R3#Sh mpls forwarding-table

Local      Outgoing    Prefix           Bytes Label   Outgoing    Next Hop
Label      Label       or Tunnel Id     Switched      interface
65300      Pop Label   2.2.2.2/32       0             Gi0/0       10.1.23.2
65301      65201       4.4.4.4/32       0             Gi0/0       10.1.23.2
65302      Pop Label   10.1.24.0/24     0             Gi0/0       10.1.23.2
```

On R2

```
RP/0/5/CPU0:R2#Sh mpls forwarding

Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------
65200  Pop         3.3.3.3/32         Gi0/1/0/3    10.1.23.3       2798
65201  Pop         4.4.4.4/32         Gi0/1/0/0    10.1.24.4       2438
```

On R4

```
RP/0/5/CPU0:R4#Sh mpls forwarding

Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------
65400  Pop         2.2.2.2/32         Gi0/1/0/0    10.1.24.2       2684
65401  65200       3.3.3.3/32         Gi0/1/0/0    10.1.24.2       0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
65402   Pop          10.1.23.0/24       Gi0/1/0/0    10.1.24.2       0
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0
```

Note the exchange of IPv4 routes between BGP neighbors are enabled by default on a this 7200 platform, which means the configured neighbors will NOT only establish a BGP session but they will also receive the advertised prefixes.

Since the "NO bgp default ipv4-unicast" is configured first (Before the neighbor commands), the local router will NOT establish a BGP peer session with any neighbor unless that neighbor is activated.

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#neighbor 3.3.3.3
RP/0/5/CPU0:R4(config-bgp-nbr)#remote-as 65001
RP/0/5/CPU0:R4(config-bgp-nbr)#update-source lo0
RP/0/5/CPU0:R4(config-bgp-nbr)#commit
```

Note the exchange of IPv4 routes between BGP neighbors is NOT enabled by default on an IOS-XR platform. The BGP address-family must be declared directly under the process to be activated. The neighbor must also be manually configured to participate in any address-family that has been declared.

To verify the configuration:

On R3

```
R3#Show ip bgp
R3#

R3#Show ip bgp summary
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

R3#

Note there is no IPv4 neighbor adjacency established between the two routers. Since the task states that these two BGP speakers should ONLY establish a VPNv4 peer session, the neighbors MUST be activated in the address-family VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#address-family vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-af)#neighbor 3.3.3.3
RP/0/5/CPU0:R4(config-bgp-nbr)#address-family vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-nbr-af)#commit
```

You should see the following console message on R3 :

```
%BGP-5-ADJCHANGE: neighbor 4.4.4.4 Up
```

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

```
RP/0/5/CPU0:R4#sh bgp vpnv4 unicast summary | b Neighbor

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down
St/PfxRcd
3.3.3.3           0 65001       4       3        1    0    0 00:01:06
0
```

NOTE: The BGP speakers have ONLY established a VPNv4 peering.

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4#Sh bgp summ | b Neigbor
```

On R4

```
RP/0/5/CPU0:R4#Sh bgp summ

% None of the requested address families are configured for instance
'default'
```

On R2

```
RP/0/5/CPU0:R2#Show bgp summ

% BGP instance 'default' not active
```

Note R2 is NOT running BGP at all, and ONLY the edge routers are running BGP. Running BGP on the core routers (P) is NOT necessary; these routers ONLY perform label switching.

Task 4

Configure a **Virtual Routing Forwarding (VRF) Instance with a name of CA** (For Customer A), a **route-distinguisher (RD) of "1:10"** and a **route-target (RT) of "1:100"** on R3. On R4, the same route targets will be used for the vrf, but the **RD should be configured to be "1:20"** and the name of the VRF should be configured as **CB**.

The "IP vrf" command creates a new VRF and enters the global configuration mode for that specific VRF. The name of the VRF is locally significant and it's case sensitive. VRF is NOT operational unless the RD is defined.

The "RD" VRF configuration mode command is used to define and assign an RD to a VRF, remember that the VRF is NOT operational without an RD. RD is a 64 bit value used to transform 32 bit customer IPv4 address, which is NON-Unique into a Unique 96 bit addresses called VPNv4. These addresses are ONLY exchanged between the PE routers and NEVER between the CE routers.

When the CE router sends an update to a PE router, the PE router prepends a 64 bit RD to the IPv4 address (32 bit address) resulting in a globally unique 96 bit address called VPNv4.

The PE router will then send the VPNv4 address/es via MP-BGP session to the other PE router/s. The receiving PE router strips the RD from the VPNv4 prefix, resulting in an IPv4 address.

Remember that the RD does NOT indicate which VRF a given prefix belongs to, it's used to make the

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

VRF prefix/es unique within the cloud. RDs DO NOT identify the VPN.

Let's configure the VRFs and the RDs:

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
```

Since VRFs are locally significant and the RDs will keep the CE routes Unique within the cloud, and the receiving PE will strip off the 64 bit RD value, how does the receiving PE know which VRF does the IP address belong to? The answer is "Route-Target".

The "Route-target import|export" command defines the "RT"; An RT is a BGP extended community that indicates which routes should be exported or imported from MP-BGP into the VRF.

Basically RTs were introduced to support identifying a site that participates in more than one VPN.

The "route-target export" command specifies an RT is to be attached to every route exported from the local VRF to MP-BGP. Whereas, the "route-target import" command specifies an RT to be used as an import filter, ONLY routes matching the RT are imported into the VRF.
This implementation allows a route to have many imported or exported RTs, all to be attached to every imported or exported route.

Let's configure RTs:

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#route-target import 1:100
R3(config-vrf)#route-target export 1:100
```

On an IOS platform, the "both" keyword can be used to replace both "import" and "export" keywords with one statement.

On R4

```
RP/0/5/CPU0:R4(config)#vrf CB
RP/0/5/CPU0:R4(config-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:100
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:100
RP/0/5/CPU0:R4(config-vrf-af)#commit
```

The VRF's are configured slightly different in IOS-XR. The Address-Family MUST be noted and the RD is not configured in this context as it is a strictly used as a BGP control Plane function.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Route-Targets are BGP extended communities that are attached to the VPNv4 addresses, and communities are NOT sent unless they are configured to be sent in IOS. In IOS-XR all communities are propagated to IBGP neighbors by default.

Let's configure R3 to send the extended communities:

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

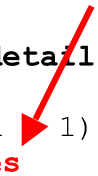
On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#address-family vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-af)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#rd 1:20
RP/0/5/CPU0:R4(config-bgp-vrf)#commit
```

Here we see where the RD is configured. It is under the VRF but the vrf is applied to the Router BGP process first.

To verify the VRF Configuration:

On R3



```
R3#Sh ip vrf detail

VRF CA (VRF Id = 1); default RD 1:10; default VPNID <not set>
  No interfaces
VRF Table ID = 1
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Think of VRFs as VLANs; once a VLAN is configured, the appropriate interface/s are assigned to that VLAN. The same philosophy applies to VRFs, you can think of VRFs as layer three VLANs, in this case, the VRFs are configured, and now the appropriate interface/s needs to be assigned to that VRF:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To associate an interface/s to a given VRF:

On R3

```
R3(config)#int s1/0
R3(config-if)#ip vrf forwarding CA
```

The above command associates an interface with the specified VRF; remember when this command is applied to a given interface, the IP address of that interface is removed and it should be reconfigured.

You should get the following console message:

```
% Interface Serial0/1 IP address 10.1.13.3 removed due to enabling VRF CA

R3(config-if)#ip address 10.1.13.3 255.255.255.0
```

Let's configure R4:

On R4

```
RP/0/5/CPU0:R4(config)#int g0/1/0/2
RP/0/5/CPU0:R4(config-if)#vrf CB
RP/0/5/CPU0:R4(config-if)#commit

% Failed to commit one or more configuration items during a pseudo-atomic
operation. All changes made have been reverted. Please issue 'show
configuration failed' from this session to view the errors

RP/0/5/CPU0:R4(config-if)#root
RP/0/5/CPU0:R4(config)#Show config failed

!! SEMANTIC ERRORS: This configuration was rejected by
!! the system due to semantic errors. The individual
!! errors with each failed configuration command can be
!! found below.

interface GigabitEthernet0/1/0/2
 vrf CB
!!% 'RSI' detected the 'fatal' condition 'The interface's numbered and
unnumbered IPv4/IPv6 addresses must be removed prior to changing or
deleting the VRF'
!
end
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config)#clear
RP/0/5/CPU0:R4(config)#int g0/1/0/2
RP/0/5/CPU0:R4(config-if)#No ipv4 address
RP/0/5/CPU0:R4(config-if)#vrf CB
RP/0/5/CPU0:R4(config-if)#ipv4 address 10.1.45.4/24
RP/0/5/CPU0:R4(config-if)#commit
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Show vrf CB detail

VRF CB; RD 1:20; VPN ID not set
VRF mode: Regular
Description not set
Interfaces:
  GigabitEthernet0/1/0/2
Address family IPV4 Unicast
  Import VPN route-target communities:
    RT:1:100
  Export VPN route-target communities:
    RT:1:100
  No import route policy
  No export route policy
Address family IPV6 Unicast
  No import VPN route-target communities
  No export VPN route-target communities
  No import route policy
  No export route policy
```

The above output begins with the name of the VRF, the default Route Distinguisher value and the VPNID which is normally not set by default. The VPNID is an extension used to further identify the VPN by using the customer "OUI" and "vpn index" that can be keyed in as a Hexadecimal or Decimal formatted number.

The interfaces that have the "VRF CB" applied will be listed under the Interfaces: list along with the warning that Connected addresses are no longer in the global routing table. It is a good practice to have at least 1 route target that the router will be both EXPORTING and IMPORTING. The 1:100 Route Target follows this basic design rule.

Currently there is no Export or Import route-maps applied. VRF-SELECT will be discussed later.

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Show ip vrf CA

Name                              Default RD          Interfaces
CA                                1:10                Se0/1

R3#Sh ip vrf interfaces

Interface              IP-Address      VRF                            Protocol
Se1/0                  10.1.13.3       CA                             up
```

To verify the connectivity between PEs and the CEs:

On R4:

```
R4#Ping 10.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.45.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

A regular ping will no longer work. A ping with no other key words will default to using the global routing table. The 10.1.45.0 /24 prefix is not accessible in the global routing table anymore. The ping must be added with the proper VRF keyword. Remember the IP address is in the VRF and NOT the global routing table.

```
R4#Ping VRF CB 10.1.45.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.45.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/32 ms
```

On R3:

```
R3#Ping VRF CA 10.1.13.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.1.13.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/29 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Configure a static default route on each Customer router located in CA and CB; these static routes should be configured to point to their respective PE router (**R3** for R1 and **R4** for R5). The PE Routers (R3 and R4), should both be configured with a static route that reaches the loopback and F0/0 interface of the Customer router, R3 and R4 should be able to see both static routes in their BGP table.

To configure a default route on the CEs:

On R1

```
R1(config)#ip route 0.0.0.0 0.0.0.0 10.1.13.3
```

On R5

```
R5(config)#ip route 0.0.0.0 0.0.0.0 10.1.45.4
```

Just a regular default route from the customer's perspective. This method is one of the recommended methods that most Service Providers prefer when offering MPLS VPN as a service to a very small company.

To configure a static route on the PEs:

Note the output of the following show command reveals that when the VRF forwarding was enabled on S1/0 interface, the "Address-family IPv4 VRF CA" was added to the BGP configuration.

On R4

```
RP/0/5/CPU0:R4(config)#router static
RP/0/5/CPU0:R4(config-static)#vrf CB
RP/0/5/CPU0:R4(config-static-vrf)#address-family ipv4 uni
RP/0/5/CPU0:R4(config-static-vrf-afi)#5.0.0.0/8 10.1.45.5
RP/0/5/CPU0:R4(config-static-vrf-afi)#192.168.5.0/24 10.1.45.5
RP/0/5/CPU0:R4(config-static-vrf-afi)#commit
```

On R3

```
R3(config)#ip route vrf CA 1.0.0.0 255.0.0.0 10.1.13.1
R3(config)#ip route vrf CA 192.168.1.0 255.255.255.0 10.1.13.1
```

To verify and test the configuration:

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Show ip route vrf CA | b Gateway
Gateway of last resort is not set

S     1.0.0.0/8 [1/0] via 10.1.13.1
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.1.13.0/24 is directly connected, Serial1/0
L        10.1.13.3/32 is directly connected, Serial1/0
S     192.168.1.0/24 [1/0] via 10.1.13.1

R3#Ping vrf CA 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms

R3#Ping vrf CA 192.168.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/28/28 ms
```

On R4

```
RP/0/5/CPU0:R4#Show route vrf CB | b Gateway
Gateway of last resort is not set

S    5.0.0.0/8 [1/0] via 10.1.45.5, 00:07:41
C    10.1.45.0/24 is directly connected, 00:27:56, GigabitEthernet0/1/0/2
L    10.1.45.4/32 is directly connected, 00:27:56, GigabitEthernet0/1/0/2
S    192.168.5.0/24 [1/0] via 10.1.45.5, 00:07:41

R4#Ping vrf CB 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 28/29/32 ms

R4#Ping vrf CB 192.168.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

In the last step of this task, the static routes that were configured on the PE routers should be redistributed into BGP so the neighboring PE router can see the route in its BGP table. This redistribution should be done under another address-family called "IPv4 Unicast"; this address-family is created automatically when the VRF, RD and the RTs are configured; the following show command verifies this fact:

On R4

```
R4#Sh run | B router bgp

RP/0/5/CPU0:R4#sh run router bgp
Sat Apr 14 22:59:03.553 UTC
router bgp 65001
 address-family vpnv4 unicast
 !
 neighbor 3.3.3.3
  remote-as 65001
  update-source Loopback0
  address-family vpnv4 unicast
  !
 !
 vrf CB
  rd 1:20
```

On R3

```
R3#Sh run | B router bgp

router bgp 65001
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 4.4.4.4 remote-as 65001
 neighbor 4.4.4.4 update-source Loopback0
 !
 address-family vpnv4
  neighbor 4.4.4.4 activate
  neighbor 4.4.4.4 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf CA
  no synchronization
 exit-address-family
```

The address-family shows up automatically in most IOS platforms but may not show up on others. If it does NOT show up, the next step will configure the address-family to redistribute the static routes. To configure the redistribution of the static routes:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#address-family vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-af)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-bgp-vrf-af)#redistribute static
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
RP/0/5/CPU0:R4(config-bgp-vrf-af)#end
```

To verify the configuration:

On R3

```
R3#Show ip route vrf CA | b Gateway
Gateway of last resort is not set

S    1.0.0.0/8 [1/0] via 10.1.13.1
B    5.0.0.0/8 [200/0] via 4.4.4.4, 00:00:27
B    192.168.5.0/24 [200/0] via 4.4.4.4, 00:00:27
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
S    192.168.1.0/24 [1/0] via 10.1.13.1
```

NOTE: The redistributed routes show up on R3 as BGP routes. Let's configure R3 to redistribute its static routes:

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#Address-family IPv4 VRF CA
R3(config-router-af)#Redistribute static
```

On R4

```
R4#Show ip route vrf CB | B Gateway
Gateway of last resort is not set

B    1.0.0.0/8 [200/0] via 3.3.3.3, 00:02:51
S    5.0.0.0/8 [1/0] via 10.1.45.5
S    192.168.5.0/24 [1/0] via 10.1.45.5
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.45.0 is directly connected, Serial0/1
B    192.168.1.0/24 [200/0] via 3.3.3.3, 00:02:51
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Note the redistributed routes should be verified on R3, if the configuration is performed successfully, R3 should be able to see networks 5.0.0.0 /8 and 192.168.5.0 /24 in its vrf CA routing table.

To verify the routes in VPNv4:

On R3

```
R3#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                0         32768 ?
*>i5.0.0.0          4.4.4.4                  0    100      0 ?
*> 192.168.1.0      10.1.13.1                0         32768 ?
*>i192.168.5.0      4.4.4.4                  0    100      0 ?
Route Distinguisher: 1:20
*>i5.0.0.0          4.4.4.4                  0    100      0 ?
*>i192.168.5.0      4.4.4.4                  0    100      0 ?
```

On R4

```
R4#Show ip bgp vpnv4 unicast | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10
*>i1.0.0.0          3.3.3.3                  0    100      0 ?
*>i192.168.1.0      3.3.3.3                  0    100      0 ?
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0          3.3.3.3                  0    100      0 ?
*> 5.0.0.0          10.1.45.5                0         32768 ?
*>i192.168.1.0      3.3.3.3                  0    100      0 ?
*> 192.168.5.0      10.1.45.5                0         32768 ?
```

One of the reasons to use different Route Distinguishers for ALL VRFs in the network is for clarity, especially when trying to figure out which VPN a route belongs to and where it's originated. As you can see, the "1:20" RD is local, but the 1.0.0.0 & 192.168.1.0 networks are present in both RDs, which means that 1.0.0.0 & 192.168.1.0 networks are participating in the same VPN. Even though the VPN is defined by the Route Target, it is still apparent that the prefixes are being imported and exported into both VRFs.

To test the configuration:

On R1

```
R1#Ping 5.5.5.5 source 1.1.1.1
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/60 ms

R1#Ping 5.5.5.5 source 192.168.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 192.168.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 84/84/88 ms
```

On R5

```
R5#Ping 1.1.1.1 source 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 5.5.5.5
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/57 ms

R5#Ping 1.1.1.1 source 192.168.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 192.168.5.5
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 84/84/84 ms
```

Note the source of the two ping commands are specified to be the loopback 0 & F0/0's IP address, if the source is NOT specified, the Ping will NOT be successful:

On R5

```
R5#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The ping failed because the source was NOT set based on the loopback0 or the F0/0 interface, if the source is not specified, the source IP address will be the closest interface, in this case the IP address of the S0/1 interface. In order to provide the capability of Pinging without specifying the source IP address, the PE routers should also redistribute the IP address of their S0/1 interface, this can be done using the "Redistribute Connected":

On R3

```
R3(config)#router bgp 65001
R3(config-router)#Address-family IPv4 vrf CA
R3(config-router-af)#Redistribute connected
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Sh route vrf CB | b Gateway
Gateway of last resort is not set

B    1.0.0.0/8 [200/0] via 3.3.3.3 (nexthop in vrf default), 00:13:58
S    5.0.0.0/8 [1/0] via 10.1.45.5, 00:22:38
B    10.1.13.0/24 [200/0] via 3.3.3.3 (nexthop in vrf default), 00:00:15
C    10.1.45.0/24 is directly connected, 00:22:38, GigabitEthernet0/1/0/2
L    10.1.45.4/32 is directly connected, 00:22:38, GigabitEthernet0/1/0/2
B    192.168.1.0/24 [200/0] via 3.3.3.3 (nexthop in vrf default),
00:13:58
S    192.168.5.0/24 [1/0] via 10.1.45.5, 00:22:38
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#address vpnv4 uni
RP/0/5/CPU0:R4(config-bgp-af)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#address ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-af)#redistribute connected
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
RP/0/5/CPU0:R4(config-bgp-vrf-af)#end
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R3

```
R3#Sh ip route vrf CA | b Gateway
Gateway of last resort is not set

S     1.0.0.0/8 [1/0] via 10.1.13.1
B     5.0.0.0/8 [200/0] via 4.4.4.4, 00:17:08
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C        10.1.13.0/24 is directly connected, Serial1/0
L        10.1.13.3/32 is directly connected, Serial1/0
B        10.1.45.0/24 [200/0] via 4.4.4.4, 00:01:12
S     192.168.1.0/24 [1/0] via 10.1.13.1
B     192.168.5.0/24 [200/0] via 4.4.4.4, 00:17:08
```

To test the connectivity:

On R1

```
R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

On R5

```
R5#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/57 ms
```

To verify the extended community attached to a given prefix:

On R4

```
RP/0/5/CPU0:R4#Sh bgp vpnv4 uni vrf CB 5.0.0.0

BGP routing table entry for 5.0.0.0/8, Route Distinguisher: 1:20
Versions:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Process            bRIB/RIB  SendTblVer
  Speaker                   4           4
    Local Label: 65403
Last Modified: Apr 14 23:16:20.621 for 00:32:04
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    3.3.3.3
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    3.3.3.3
  Local
    10.1.45.5 from 0.0.0.0 (4.4.4.4)
      Origin incomplete, metric 0, localpref 100, weight 32768, valid,
redistributed, best, group-best, import-candidate
      Received Path ID 0, Local Path ID 1, version 4
      Extended community: RT:1:100
```

We have successfully verified that the route-target has been added to the routes as they are exported from the local VRF. This is important as NO vrf will be able to import this route if the extended community is not added to the prefix. The VPNv4 label has also been added to the prefix. The "65403(Local Label) and the nolabel (Rcvd)" shows that this router expects to see this route with label 65403 when it recieves the data traffic and will remove the label before forwarding it to the Customer (R5) router.

```
RP/0/5/CPU0:R4#Sh bgp vpnv4 uni labels | b Netw

Sat Apr 14 23:50:11.534 UTC
   Network            Next Hop        Rcvd Label      Local Label
Route Distinguisher: 1:10
*>i1.0.0.0/8          3.3.3.3         65303           nolabel
*>i10.1.13.0/24       3.3.3.3         65305           nolabel
*>i192.168.1.0/24     3.3.3.3         65304           nolabel
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0/8          3.3.3.3         65303           nolabel
*> 5.0.0.0/8          10.1.45.5       nolabel         65403
*>i10.1.13.0/24       3.3.3.3         65305           nolabel
*> 10.1.45.0/24       0.0.0.0         nolabel         65405
*>i192.168.1.0/24     3.3.3.3         65304           nolabel
*> 192.168.5.0/24     10.1.45.5       nolabel         65404
```

This command displays the VPNv4 labels that are added to the routes. A common question that is often asked is "What does the aggregate keyword mean?" when it shows up in the table. The aggregate keyword means that the route requires an IP lookup to determine the next hop of this packet. When we redistributed the connected routes, the far end does not know about this next hop due to the fact that this route is not in the VPN.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

By redistributing the connected route, a normal ping from the interface closest to the destination can be performed by R1 due to the fact that R5 now knows about the source of the packet.

```
RP/0/5/CPU0:R4#Sh mpls forwarding vrf CB

Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------
------
65403  Unlabelled  5.0.0.0/8[V]       Gi0/1/0/2    10.1.45.5       0
65404  Unlabelled  192.168.5.0/24[V]  Gi0/1/0/2    10.1.45.5       0
65405  Aggregate   CB: Per-VRF Aggr[V]    \
                                      CB                           0
```

On R3

```
R3#Show mpls forwarding-table vrf CA

Local  Outgoing    Prefix             Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id       switched   interface
303    Untagged    1.0.0.0/8[V]       2600       Se0/1      point2point
304    Untagged    192.168.1.0/24[V]  520        Se0/1      point2point
305    Aggregate   10.1.13.0/24[V]    520
```

Task 6

Remove the Static routes and replace the current method of routing between the PE and Customers with RIPv2 Routing.

Removing the Static Configuration on the CE routers:

On R1

```
R1(config)#NO ip route 0.0.0.0 0.0.0.0
```

On R5

```
R5(config)#NO ip route 0.0.0.0 0.0.0.0
```

Removing the Static Configuration on the PE routers:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#NO ip route vrf CA 1.0.0.0 255.0.0.0

R3(config)#NO ip route vrf CA 192.168.1.0 255.255.255.0

R3(config)#router bgp 65001
R3(config-router)# address ipv4 vrf CA
R3(config-router-af)#NO redistribute static
R3(config-router-af)#NO redistribute connected
```

On R4

```
RP/0/5/CPU0:R4(config)#NO router static
RP/0/5/CPU0:R4(config)#commit

RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#address-family ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-af)#NO redistribute static
RP/0/5/CPU0:R4(config-bgp-vrf-af)#NO redistribute connected
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
```

Configuring RIPv2 routing on CE routers:

On R1

```
R1(config)#router rip
R1(config-router)#No au
R1(config-router)#ver 2
R1(config-router)#netw 10.0.0.0
R1(config-router)#netw 192.168.1.0
R1(config-router)#netw 1.0.0.0
```

On R5

```
R5(config)#router rip
R5(config-router)#No au
R5(config-router)#ver 2
R5(config-router)#network 10.0.0.0

R5(config-router)#netw 5.0.0.0
R5(config-router)#netw 192.168.5.0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Configuring RIPv2 routing on PE routers:

On R3

```
R3(config)#router rip
R3(config-router)#ver 2
R3(config-router)#Address-family ipv4 vrf CA

R3(config-router-af)#Network 10.0.0.0
R3(config-router-af)#version 2
R3(config-router-af)#No au
```

On R4

```
RP/0/5/CPU0:R4(config)#router rip
RP/0/5/CPU0:R4(config-rip)#vrf CB
RP/0/5/CPU0:R4(config-rip-vrf)#No auto
RP/0/5/CPU0:R4(config-rip-vrf)#int g0/1/0/2
```

To verify the configuration:

On R3

```
R3#Sh ip route vrf CA | b Gateway
Gateway of last resort is not set

      1.0.0.0/32 is subnetted, 1 subnets
R        1.1.1.1 [120/1] via 10.1.13.1, 00:00:06, Serial1/0
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.1.13.0/24 is directly connected, Serial1/0
L        10.1.13.3/32 is directly connected, Serial1/0
R     192.168.1.0/24 [120/1] via 10.1.13.1, 00:00:06, Serial1/0
```

On R4

```
RP/0/5/CPU0:R4#Sh route vrf CB | b Gateway
Gateway of last resort is not set

R    5.5.5.5/32 [120/1] via 10.1.45.5, 00:01:46, GigabitEthernet0/1/0/2
C    10.1.45.0/24 is directly connected, 17:53:07, GigabitEthernet0/1/0/2
L    10.1.45.4/32 is directly connected, 17:53:07, GigabitEthernet0/1/0/2
R    192.168.5.0/24 [120/1] via 10.1.45.5, 00:01:46,
GigabitEthernet0/1/0/2
```

Note the routes are in the appropriate VRFs, the next step is to redistribute the routes.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The redistribution of RIPv2 into MP-BGP is necessary for the routes to show up on the other PE Router.

The redistribution of MP-BGP into RIP is necessary for the local router to translate the routes that have been received by the remote PE router so the CE router/s can see the routes.

Step One

RIP routes are redistributed into the BGP for the specified VRF:

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-bgp-vrf-af)#redistribute rip
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
```

To verify the configuration:

Because of the redistribution, the PE router (R3) on the other side can now see the routes in its BGP table:

On R3

```
R3#Sh ip bgp vpnv4 vrf CA | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*>i5.5.5.5/32       4.4.4.4                  1    100      0 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
*>i192.168.5.0      4.4.4.4                  1    100      0 ?
```

Step Two

In this step the routes are redistributed into RIP, so the CE router will have them in its routing table:

On R3

When the routes are redistributed into RIP, a metric of 5 is assigned; this is done on purpose so they can easily be identified in the routing table of the Customer router (R1):

```
R3(config)#router rip
R3(config-router)#Address-family ipv4 vrf CA
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router-af)#redistribute BGP 65001 metric 5
```

To verify the configuration:

On R1

Note R1 (The CE router) has the routes in its routing table:

```
R1#Show ip route rip | Inc R

R    5.0.0.0/8 [120/5] via 10.1.13.3, 00:00:08, Serial0/1
R    192.168.5.0/24 [120/5] via 10.1.13.3, 00:00:08, Serial0/1
R       10.1.45.0 [120/5] via 10.1.13.3, 00:00:08, Serial0/1
```

The same needs to be done in reverse order from R3 to R4 and verified on R5:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#Address-family ipv4 vrf CA
R3(config-router-af)#Redistribute RIP
```

On R4

```
RP/0/5/CPU0:R4(config)#router rip
RP/0/5/CPU0:R4(config-rip)#vrf CB
RP/0/5/CPU0:R4(config-rip-vrf)#redistribute bgp 65001 internal
RP/0/5/CPU0:R4(config-rip-vrf)#commit
```

To verify the configuration:

On R5

```
R5#Show ip route rip | Inc R

R5#sh ip route rip | i R
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route, H - NHRP
R        1.1.1.1 [120/2] via 10.1.45.4, 00:00:01, GigabitEthernet0/0
R     192.168.1.0/24 [120/2] via 10.1.45.4, 00:00:01, GigabitEthernet0/0
```

To see the full picture, the following show commands verifies network 1.0.0.0 /8 from R1 all the way to the upstream router R5; this is called the control plane.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Sh ip route 1.0.0.0

Routing entry for 1.0.0.0/8, 2 known subnets
  Attached (2 connections)
  Variably subnetted with 2 masks
  Redistributing via rip
C        1.0.0.0/8 is directly connected, Loopback0
L        1.1.1.1/32 is directly connected, Loopback0

R1#Sh ip route 1.1.1.1

Routing entry for 1.1.1.1/32
  Known via "connected", distance 0, metric 0 (connected)
  Routing Descriptor Blocks:
  * directly connected, via Loopback0
      Route metric is 0, traffic share count is 1

R1#Show ip rip database 1.0.0.0 255.0.0.0

1.0.0.0/8    directly connected, Loopback0
```

On R3

```
R3#Sh ip route vrf CA | b Gateway
Gateway of last resort is not set

      1.0.0.0/8 is subnetted, 1 subnets
R        1.0.0.0 [120/1] via 10.1.13.1, 00:00:30, Serial1/0
      5.0.0.0/8 is subnetted, 1 subnets
B        5.0.0.0 [200/1] via 4.4.4.4, 00:03:37
      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
C        10.1.13.0/24 is directly connected, Serial1/0
L        10.1.13.3/32 is directly connected, Serial1/0
B        10.1.45.0/24 [200/0] via 4.4.4.4, 00:18:50
R     192.168.1.0/24 [120/1] via 10.1.13.1, 00:00:30, Serial1/0
B     192.168.5.0/24 [200/1] via 4.4.4.4, 00:18:50

R3#Sh ip route vrf CA 1.1.1.1

Routing Table: CA
Routing entry for 1.0.0.0/8
  Known via "rip", distance 120, metric 1
  Redistributing via bgp 65001, rip
  Advertised by bgp 65001
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Last update from 10.1.13.1 on Serial1/0, 00:00:11 ago
  Routing Descriptor Blocks:
  * 10.1.13.1, from 10.1.13.1, 00:00:11 ago, via Serial1/0
      Route metric is 1, traffic share count is 1

R3#Sh ip bgp vpnv4 vrf CA | b Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                1         32768 ?
*>i5.0.0.0          4.4.4.4                  1    100      0 ?
*> 10.1.13.0/24     0.0.0.0                  0         32768 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
*> 192.168.1.0      10.1.13.1                1         32768 ?
*>i192.168.5.0      4.4.4.4                  1    100      0 ?

R3#Sh ip bgp vpnv4 vrf CA 1.0.0.0/8

BGP routing table entry for 1:10:1.0.0.0/8, version 29
Paths: (1 available, best #1, table CA)
  Advertised to update-groups:
     1
  Local
    10.1.13.1 from 0.0.0.0 (3.3.3.3)
      Origin incomplete, metric 1, localpref 100, weight 32768, valid,
sourced, best
      Extended Community: RT:1:100
      mpls labels in/out 65306/nolabel
```

Note the Origin code is (incomplete) because the route was redistributed and therefore, the origin of the route is unknown. The Weight attribute is 32768, because the local router is advertising the route in BGP. The extended community is identified, this is the configured route-target, and the last line states that if the local router receives a packet with a label of 65306 it will remove the label and forward it.

On R4

```
RP/0/5/CPU0:R4#Sh bgp vpnv4 unicast vrf CB 1.0.0.0

BGP routing table entry for 1.0.0.0/8, Route Distinguisher: 1:20
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 33          33
Last Modified: Apr 15 17:20:03.621 for 00:09:38
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Not advertised to any peer
  Local
    3.3.3.3 (metric 3) from 3.3.3.3 (3.3.3.3)
      Received Label 65306
      Origin incomplete, metric 1, localpref 100, valid, internal, best,
group-best, import-candidate, imported
      Received Path ID 0, Local Path ID 1, version 33
      Extended community: RT:1:100

RP/0/5/CPU0:R4#Sh bgp vpnv4 unicast vrf CB | b Network

   Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0/8          3.3.3.3                  1    100      0 ?
*> 5.0.0.0/8          10.1.45.5                1         32768 ?
*>i10.1.13.0/24       3.3.3.3                  0    100      0 ?
*> 10.1.45.0/24       0.0.0.0                  0         32768 ?
*>i192.168.1.0/24     3.3.3.3                  1    100      0 ?
*> 192.168.5.0/24     10.1.45.5                1         32768 ?

RP/0/5/CPU0:R4#Sh route vrf CB | b Gateway
Gateway of last resort is not set

B    1.0.0.0/8 [200/1] via 3.3.3.3 (nexthop in vrf default), 00:12:19
R    5.0.0.0/8 [120/1] via 10.1.45.5, 00:11:59, GigabitEthernet0/1/0/2
B    10.1.13.0/24 [200/0] via 3.3.3.3 (nexthop in vrf default), 00:21:58
C    10.1.45.0/24 is directly connected, 18:24:37, GigabitEthernet0/1/0/2
L    10.1.45.4/32 is directly connected, 18:24:37, GigabitEthernet0/1/0/2
B    192.168.1.0/24 [200/1] via 3.3.3.3 (nexthop in vrf default), 00:21:58
R    192.168.5.0/24 [120/1] via 10.1.45.5, 00:33:16, GigabitEthernet0/1/0/2

RP/0/5/CPU0:R4#Sh route vrf CB 1.0.0.0

Routing entry for 1.0.0.0/8
  Known via "bgp 65001", distance 200, metric 1, type internal
  Installed Apr 15 17:20:03.584 for 00:13:04
  Routing Descriptor Blocks
    3.3.3.3, from 3.3.3.3
      Nexthop in Vrf: "default", Table: "default", IPv4 Unicast, Table
Id: 0xe0000000
      Route metric is 1
  No advertising protos.
```

On R5

```
R5#Sh ip route rip | i R
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route, H - NHRP
R        1.0.0.0 [120/2] via 10.1.45.4, 00:00:22, GigabitEthernet0/0
R     192.168.1.0/24 [120/2] via 10.1.45.4, 00:00:22, GigabitEthernet0/0
```

To test the configuration:

On R5

```
R5#Ping 1.1.1.1 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
Packet sent with a source address of 5.5.5.5
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R1

```
R1#Ping 5.5.5.5 so lo0

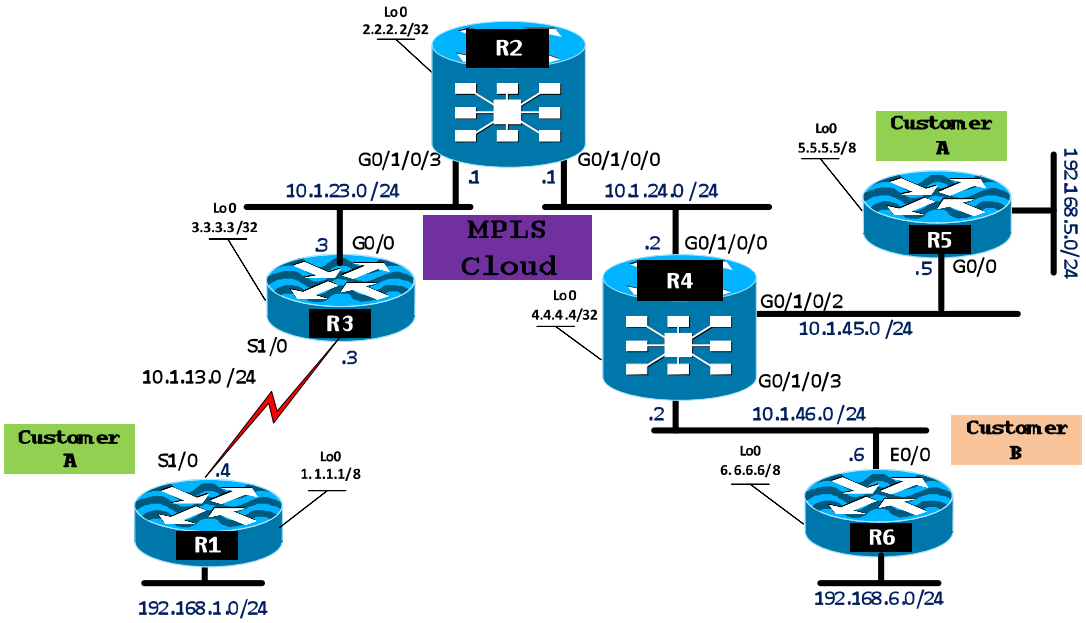
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
Packet sent with a source address of 1.1.1.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 3
OSPF Routing in a VPN



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100, R5 in VLAN 500 and F0/1 interface of R6 in VLAN 600.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the F0/0 interface of R4 and R6 in VLAN 900.
- Configure the rest of the routers according to the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

On R2

```
RP/0/5/CPU0:R2(config)#router ospf TST
RP/0/5/CPU0:R2(config-ospf)#area 0
RP/0/5/CPU0:R2(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R2(config-ospf-ar-if)#passive
RP/0/5/CPU0:R2(config-ospf-ar-if)#commit
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf TST
RP/0/5/CPU0:R4(config-ospf)#area 0
RP/0/5/CPU0:R4(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R4(config-ospf-ar-if)#passive
RP/0/5/CPU0:R4(config-ospf-ar-if)#commit
```

To verify the configuration:

On R3

```
R3#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O        2.2.2.2 [110/2] via 10.1.23.2, 03:36:56, GigabitEthernet0/0
O        4.4.4.4 [110/3] via 10.1.23.2, 03:35:40, GigabitEthernet0/0
O        10.1.24.0/24 [110/2] via 10.1.23.2, 03:36:56, GigabitEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
RP/0/5/CPU0:R2#sh route ospf | i O

O    3.3.3.3/32 [110/2] via 10.1.23.3, 03:37:56, GigabitEthernet0/1/0/3
O    4.4.4.4/32 [110/2] via 10.1.24.4, 03:36:40, GigabitEthernet0/1/0/0
RP/0/5/CPU0:R2#
```

On R4

```
RP/0/5/CPU0:R4#Sh route ospf | i O

O    2.2.2.2/32 [110/2] via 10.1.24.2, 03:37:22, GigabitEthernet0/1/0/0
O    3.3.3.3/32 [110/3] via 10.1.24.2, 03:37:22, GigabitEthernet0/1/0/0
O    10.1.23.0/24 [110/2] via 10.1.24.2, 03:37:22, GigabitEthernet0/1/0/0
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 65200 – 65299
R3 – 65300 – 65399
R4 – 65400 – 65499

On R3

```
R3(config)#mpls label range 65300 65399

R3(config)#mpls label prot ldp

R3(config)#mpls ldp router-id lo0

R3(config)#int g0/0
R3(config-if)#MPLS IP
```

On R2

```
RP/0/5/CPU0:R2(config)#mpls label range 65200 65299
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R2(config)#mpls ldp
RP/0/5/CPU0:R2(config-ldp)#router-id 2.2.2.2
RP/0/5/CPU0:R2(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ldp-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ldp-if)#commit
RP/0/5/CPU0:R2(config-ldp-if)#end
```

On R4

```
RP/0/5/CPU0:R4(config)#mpls label range 65400 65499

RP/0/5/CPU0:R4(config)#mpls ldp
RP/0/5/CPU0:R4(config-ldp)#router-id 4.4.4.4
RP/0/5/CPU0:R4(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ldp-if)#commit
RP/0/5/CPU0:R4(config-ldp-if)#end
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Sh mpls ldp neighbor

Peer LDP Identifier: 2.2.2.2:0
  TCP connection: 2.2.2.2:646 - 4.4.4.4:51796
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 1560/1562; Downstream-Unsolicited
  Up time: 22:40:58
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    2.2.2.2          10.1.23.2        10.1.24.2
```

On R2

```
RP/0/5/CPU0:R2#Sh mpls ldp neighbor

Peer LDP Identifier: 3.3.3.3:0
  TCP connection: 3.3.3.3:47318 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 1571/1572; Downstream-Unsolicited
  Up time: 22:45:08
  LDP Discovery Sources:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    GigabitEthernet0/1/0/3
  Addresses bound to this peer:
    3.3.3.3          10.1.13.3        10.1.23.3

Peer LDP Identifier: 4.4.4.4:0
  TCP connection: 4.4.4.4:51796 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 1565/1563; Downstream-Unsolicited
  Up time: 22:43:44
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    4.4.4.4          10.1.24.4        10.1.45.4        10.1.46.4
```

On R3

```
R3#Sh mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.47318
        State: Oper; Msgs sent/rcvd: 1574/1573; Downstream
        Up time: 22:47:30
        LDP discovery sources:
          GigabitEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          2.2.2.2         10.1.24.2       10.1.23.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 send-community extended
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#address vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-af)#neighbor 3.3.3.3
RP/0/5/CPU0:R4(config-bgp-nbr)#remote-as 65001

RP/0/5/CPU0:R4(config-bgp-nbr)#update-source loopback 0
RP/0/5/CPU0:R4(config-bgp-nbr)#address-family vpnv4 uni
RP/0/5/CPU0:R4(config-bgp-nbr-af)#commit
RP/0/5/CPU0:R4(config-bgp-nbr-af)#end
```

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

```
RP/0/5/CPU0:R4#Sh bgp VPNv4 unicast summ | b Neighbor

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
3.3.3.3           0 65001       6       4        1    0    0 00:01:45          0
```

Task 4

Configure VRFs on R3 and R4 and enable VRF forwarding on the interfaces of these two routers based on the following chart.

Router	VRF	RD	RT	Interface
R3	CA	1:10	1:105	S1/0
R4	CA	1:50	1:105	G0/1/0/2
	CB	1:60	1:600	G0/1/0/3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
R3(config-vrf)#route-target both 1:105

R3(config-vrf)#int S1/0
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip addr 10.1.13.3 255.255.255.0
```

On R4

```
RP/0/5/CPU0:R4(config)#vrf CA
RP/0/5/CPU0:R4(config-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:105
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:105
RP/0/5/CPU0:R4(config-vrf-af)#commit
RP/0/5/CPU0:R4(config-vrf-af)#exit

RP/0/5/CPU0:R4(config)#vrf CB
RP/0/5/CPU0:R4(config-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:600
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:600
RP/0/5/CPU0:R4(config-vrf-af)#commit

RP/0/5/CPU0:R4(config)#int g0/1/0/2
RP/0/5/CPU0:R4(config-if)#No ipv4 address
RP/0/5/CPU0:R4(config-if)#vrf CA
RP/0/5/CPU0:R4(config-if)#ipv4 address 10.1.45.4/24
RP/0/5/CPU0:R4(config-if)#commit

RP/0/5/CPU0:R4(config-if)#int g0/1/0/3
RP/0/5/CPU0:R4(config-if)#No ipv4 address
RP/0/5/CPU0:R4(config-if)#vrf CB
RP/0/5/CPU0:R4(config-if)#ipv4 address 10.1.46.4/24
RP/0/5/CPU0:R4(config-if)#commit
RP/0/5/CPU0:R4(config-if)#end
```

To verify the configuration:

On R3

```
R3#Sh ip vrf det

VRF CA (VRF Id = 1); default RD 1:10; default VPNID <not set>
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Interfaces:
    Se1/0
VRF Table ID = 1
  Export VPN route-target communities
    RT:1:105
  Import VPN route-target communities
    RT:1:105
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

On R4

```
RP/0/5/CPU0:R4#Sh vrf CA detail

VRF CA; RD not set; VPN ID not set
VRF mode: Regular
Description not set
Interfaces:
  GigabitEthernet0/1/0/2
Address family IPV4 Unicast
  Import VPN route-target communities:
    RT:1:105
  Export VPN route-target communities:
    RT:1:105
  No import route policy
  No export route policy
Address family IPV6 Unicast
  No import VPN route-target communities
  No export VPN route-target communities
  No import route policy
  No export route policy
```

Remember, the RD is a BGP parameter and is set when the vrf is placed under the BGP process in IOS-XR.

```
RP/0/5/CPU0:R4#Sh vrf CB detail

VRF CB; RD not set; VPN ID not set
VRF mode: Regular
Description not set
Interfaces:
  GigabitEthernet0/1/0/3
Address family IPV4 Unicast
  Import VPN route-target communities:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    RT:1:600
  Export VPN route-target communities:
    RT:1:600
  No import route policy
  No export route policy
Address family IPV6 Unicast
  No import VPN route-target communities
  No export VPN route-target communities
  No import route policy
  No export route policy  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Task 5

Configure Customers R1 and R5 with a VRF service that incorporates OSPF as the routing protocol; R3 should use OSPF process-id of 3, whereas, R4 should use process-id of 4 for this purpose, R1 and R5's loopback0 & F0/0 interfaces should be advertised with their correct mask.
Configure the VRF applied to R6 to use RIPv2 and ensure that this customer participates in the same Closed User Group (CUG) as R1 and R5. R6 should advertise its Loopback0 and F0/0 interfaces in its routing domain.

To configure the CE routers:

On R1

```
R1(config)#router ospf 1

R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#Netw 10.1.13.1 0.0.0.0 area 0
R1(config-router)#Netw 192.168.1.1 0.0.0.0 area 0

R1(config-router)#Int lo0
R1(config-if)#ip ospf netw point-to-point
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#Netw 5.5.5.5 0.0.0.0 area 0
R5(config-router)#Netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#Netw 192.168.5.5 0.0.0.0 area 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R5(config-router)#int lo0
R5(config-if)#ip ospf net point-to-point
```

On R6

```
R6(config)#router rip
R6(config-router)#No au
R6(config-router)#ver 2
R6(config-router)#Netw 192.168.6.0
R6(config-router)#Netw 6.0.0.0
R6(config-router)#Netw 10.0.0.0
```

To configure the PE routers:

On R3

```
R3(config)#Router ospf 3 vrf CA
R3(config-router)#network 10.1.13.3 0.0.0.0 area 0
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 4
RP/0/5/CPU0:R4(config-ospf)#vrf CA
RP/0/5/CPU0:R4(config-ospf-vrf)#area 0
RP/0/5/CPU0:R4(config-ospf-vrf-ar)#int g0/1/0/2
RP/0/5/CPU0:R4(config-ospf-vrf-ar-if)#commit
RP/0/5/CPU0:R4(config-ospf-vrf-ar-if)#root

RP/0/5/CPU0:R4(config)#router rip
RP/0/5/CPU0:R4(config-rip)#vrf CB
RP/0/5/CPU0:R4(config-rip-vrf)#int g0/1/0/3
RP/0/5/CPU0:R4(config-rip-vrf-if)#commit
RP/0/5/CPU0:R4(config-rip-vrf-if)#end
```

To verify the configuration:

On R3

```
R3#Sh ip ospf nei

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           1   FULL/BDR        00:00:35    10.1.23.2       FastEthernet0/0
1.1.1.1           0   FULL/  -        00:00:31    10.1.13.1       Serial0/1
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

1.1.1.1 is R1 (The CE) and 2.2.2.2 is the OSPF adjacency with R2 (The P router).

```
R3#Show ip route vrf CA | B Gateway
Gateway of last resort is not set

O    1.0.0.0/8 [110/65] via 10.1.13.1, 00:06:06, Serial0/1
     10.0.0.0/24 is subnetted, 1 subnets
C       10.1.13.0 is directly connected, Serial0/1
O    192.168.1.0/24 [110/65] via 10.1.13.1, 00:06:06, Serial0/1
```

On R4

```
P/0/5/CPU0:R4#Sh ospf nei | b Neighbors

Neighbors for OSPF TST

Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2         1     FULL/DR         00:00:32    10.1.24.2
GigabitEthernet0/1/0/0
    Neighbor is up for 1d04h

Total neighbor count: 1

RP/0/5/CPU0:R4#Sh ospf vrf CA neighbor | b Neighbors

Neighbors for OSPF 4, VRF CA

Neighbor ID     Pri   State           Dead Time   Address
Interface
5.5.5.5         1     FULL/DR         00:00:32    10.1.45.5
GigabitEthernet0/1/0/2
    Neighbor is up for 00:01:53

Total neighbor count: 1

RP/0/5/CPU0:R4#Sh route vrf CA ospf | i O

O    5.0.0.0/8 [110/2] via 10.1.45.5, 00:00:32, GigabitEthernet0/1/0/2
O    192.168.5.0/24 [110/11] via 10.1.45.5, 00:09:53, GigabitEthernet0/1/0/2
```

The CE routes are redistribted into MP-BGP:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router-af)#redistribute OSPF 3
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CA
RP/0/5/CPU0:R4(config-bgp-vrf)#rd 1:50
RP/0/5/CPU0:R4(config-bgp-vrf)#address ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-af)#redistribute ospf 4
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
RP/0/5/CPU0:R4(config-bgp-vrf-af)#root

RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#rd 1:60
RP/0/5/CPU0:R4(config-bgp-vrf)#address-family ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-af)#redistribute rip
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
RP/0/5/CPU0:R4(config-bgp-vrf-af)#end
```

To verify the configuration:

On R3

```
R3#Show ip bgp vpnv4 all | B Network

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                65         32768 ?
*>i5.0.0.0          4.4.4.4                   2    100      0 ?
*> 10.1.13.0/24     0.0.0.0                   0         32768 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
*> 192.168.1.0      10.1.13.1                65         32768 ?
*>i192.168.5.0      4.4.4.4                  11    100      0 ?
Route Distinguisher: 1:50
*>i5.0.0.0          4.4.4.4                   2    100      0 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
*>i192.168.5.0      4.4.4.4                  11    100      0 ?
```

Since none of the VRFs configured on R3 have the "route-target import 1:600", R3 will NOT have the routes from R6 in its table. The following configuration resolves this problem:

On R4

```
RP/0/5/CPU0:R4(config)#vrf CB
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config-vrf)#address ipv4 unicast
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:615
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:615
RP/0/5/CPU0:R4(config-vrf-af)#commit
RP/0/5/CPU0:R4(config-vrf-af)#root
```

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#route-target both 1:615
```

To verify the configuration:

On R3

```
R3#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                65         32768 ?
*>i5.0.0.0          4.4.4.4                   2    100      0 ?
*>i6.0.0.0          4.4.4.4                   1    100      0 ?
*> 10.1.13.0/24     0.0.0.0                   0         32768 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
*>i10.1.46.0/24     4.4.4.4                   0    100      0 ?
*> 192.168.1.0      10.1.13.1                65         32768 ?
*>i192.168.5.0      4.4.4.4                  11    100      0 ?
*>i192.168.6.0      4.4.4.4                   1    100      0 ?
Route Distinguisher: 1:50
*>i5.0.0.0          4.4.4.4                   2    100      0 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
*>i192.168.5.0      4.4.4.4                  11    100      0 ?
Route Distinguisher: 1:60
*>i6.0.0.0          4.4.4.4                   1    100      0 ?
*>i10.1.46.0/24     4.4.4.4                   0    100      0 ?
*>i192.168.6.0      4.4.4.4                   1    100      0 ?
```

On R4

```
RP/0/5/CPU0:R4#Sh bgp vpnv4 unicast | b Network

   Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10
*>i1.0.0.0/8         3.3.3.3                  65    100      0 ?
*>i10.1.13.0/24      3.3.3.3                   0    100      0 ?
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*>i192.168.1.0/24     3.3.3.3                  65    100        0 ?
Route Distinguisher: 1:50 (default for vrf CA)
*>i1.0.0.0/8          3.3.3.3                  65    100        0 ?
*> 5.0.0.0/8          10.1.45.5                 2           32768 ?
*>i10.1.13.0/24       3.3.3.3                   0    100        0 ?
*> 10.1.45.0/24       0.0.0.0                   0           32768 ?
*>i192.168.1.0/24     3.3.3.3                  65    100        0 ?
*> 192.168.5.0/24     10.1.45.5                11           32768 ?
Route Distinguisher: 1:60 (default for vrf CB)
*>i1.0.0.0/8          3.3.3.3                  65    100        0 ?
*> 6.0.0.0/8          10.1.46.6                 1           32768 ?
*>i10.1.13.0/24       3.3.3.3                   0    100        0 ?
*> 10.1.46.0/24       0.0.0.0                   0           32768 ?
*>i192.168.1.0/24     3.3.3.3                  65    100        0 ?
*> 192.168.6.0/24     10.1.46.6                 1           32768 ?
```

Redistribute BGP into OSPF on R3, and into OSPF and RIP on R4 so the CE routers (R1, R5 and R6) can see the routes in their routing table:

On R3

```
R3(config)#router ospf 3 vrf CA
R3(config-router)#redistribute bgp 65001 subnets
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 4
RP/0/5/CPU0:R4(config-ospf)#vrf CA
RP/0/5/CPU0:R4(config-ospf-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-ospf-vrf)#redistribute bgp 65001
RP/0/5/CPU0:R4(config-ospf-vrf)#root

RP/0/5/CPU0:R4(config)#router rip
RP/0/5/CPU0:R4(config-rip)#vrf CB
RP/0/5/CPU0:R4(config-rip-vrf)#redistribute bgp 65001
RP/0/5/CPU0:R4(config-rip-vrf)#default-metric 5
RP/0/5/CPU0:R4(config-rip-vrf)#commit
RP/0/5/CPU0:R4(config-rip-vrf)#end
```

To verify the configuration:

On R1

```
R1#Sh ip route ospf | i O
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O E2   5.0.0.0/8 [110/2] via 10.1.13.3, 04:01:50, Serial1/0
O E2   6.0.0.0/8 [110/1] via 10.1.13.3, 04:01:50, Serial1/0
O E2      10.1.45.0/24 [110/1] via 10.1.13.3, 04:01:50, Serial1/0
O E2      10.1.46.0/24 [110/1] via 10.1.13.3, 04:01:50, Serial1/0
O E2   192.168.5.0/24 [110/11] via 10.1.13.3, 04:01:50, Serial1/0
O E2   192.168.6.0/24 [110/1] via 10.1.13.3, 01:35:35, Serial1/0
R1#
```

On R5

```
R5#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O E2   1.0.0.0/8 [110/65] via 10.1.45.4, 00:02:54, GigabitEthernet0/0
O E2      10.1.13.0/24 [110/1] via 10.1.45.4, 00:02:54, GigabitEthernet0/0
O E2   192.168.1.0/24 [110/65] via 10.1.45.4, 00:02:54, GigabitEthernet0/0
```

On R6

```
R6#Sh ip route rip | i R

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route, H - NHRP
R      1.0.0.0/8 [120/5] via 10.1.46.4, 00:00:16, Ethernet0/0
R         10.1.13.0/24 [120/5] via 10.1.46.4, 00:00:16, Ethernet0/0
R      192.168.1.0/24 [120/5] via 10.1.46.4, 00:00:16, Ethernet0/0
```

Note the routes that are in the same ospf area are showing up as External type 2 routes.

Was this because of the redistribution?
Although redistribution results in this behavior under normal circumstances, this does not follow classic OSPF default behavior. In this case both OSPF enabled routers are configured in the same area (Area 0).

To see the reason behind the behavior:

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh ip bgp vpnv4 vrf CA 1.0.0.0

BGP routing table entry for 1:10:1.0.0.0/8, version 19
Paths: (1 available, best #1, table CA)
  Advertised to update-groups:
     1
  Local
    10.1.13.1 from 0.0.0.0 (3.3.3.3)
      Origin incomplete, metric 65, localpref 100, weight 32768, valid,
sourced, best
      Extended Community: RT:1:105 RT:1:615
        OSPF DOMAIN ID:0x0005:0x000000030200 OSPF RT:0.0.0.0:2:0
        OSPF ROUTER ID:10.1.13.3:0
      mpls labels in/out 65306/nolabel
```

Every domain-id on a Cisco Router will begin with "0005", "0105" or "0205" which identifies the type of format used. With the "0005" value used the next 4 Octect are viewed as the Global Administrator field and is used as part of the domain-id. The last 2 Octets would be used as a Local Administer field and it's ignored. As you can see by the output, the Global Administrator field is set to "3" which represents the OSPF process-id in HEX. The "0200" output at the end is ignored.
The OSPF RT:0.0.0.0:2:0 value represents the Area of OSPF "0.0.0.0", the internal ospf route "2", and the last value of the last digit "0" means that the route is neither an External 1 nor an External 2 route.

On R4

```
RP/0/5/CPU0:R4#Sh bgp vpnv4 unicast vrf CA 5.0.0.0/8

BGP routing table entry for 5.0.0.0/8, Route Distinguisher: 1:50
Versions:
  Process           bRIB/RIB  SendTblVer
  Speaker                 63          63
    Local Label: 65403
Last Modified: Apr 21 05:34:54.630 for 00:21:16
Paths: (1 available, best #1)
  Advertised to peers (in unique update groups):
    3.3.3.3
  Path #1: Received by speaker 0
  Advertised to peers (in unique update groups):
    3.3.3.3
  Local
    10.1.45.5 from 0.0.0.0 (4.4.4.4)
      Origin incomplete, metric 2, localpref 100, weight 32768, valid,
redistributed, best, group-best, import-candidate
      Received Path ID 0, Local Path ID 1, version 63
      Extended community: RT:1:105 OSPF route-type:0:1:0x0 OSPF router-
id:4.4.4.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The domain-id for the prefixes does not match. **IN IOS-XR the domai-id is not set. Under normal OSPF design rules, the process-id of an OSPF router does not need to match any other peer's process-id. However, with MPLS, the router uses the OSPF process as a portion of the domain-id. Therefore, we must manually change the process-id to match on both sides of the cloud or statically configure a domain-id that matches on both routers.**

Routers that do not share the same domain-id are considered type 5 or external to OSPF. If the domain-ids match, the routes are considered type "3" LSAs.

Manually configuring the domain-id is preferred as this does not change the classic understanding of the process-id in OSPF. A clearing of BGP with the "clear ip bgp * in" and "clear ip bgp * out" will speed up the convergence.

On R3

```
R3(config)#router ospf 3 vrf CA
R3(config-router)#domain-id type 0005 value 010101010101
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 4
RP/0/5/CPU0:R4(config-ospf)#vrf CA
RP/0/5/CPU0:R4(config-ospf-vrf)#domain-id typ 0005 value 010101010101
RP/0/5/CPU0:R4(config-ospf-vrf)#commit
RP/0/5/CPU0:R4(config-ospf-vrf)#end
```

On R3

```
R3#Show ip bgp vpnv4 vrf CA 1.0.0.0 | Inc DOMAIN

       OSPF DOMAIN ID:0x0005:0x000000010200 OSPF RT:0.0.0.0:2:0

R3#sh ip bgp vpnv4 vrf CA 1.0.0.0 | inc DOMAIN
        OSPF DOMAIN ID:0x0005:0x010101010101 OSPF RT:0.0.0.0:2:0
```

On R4

Domain-ids match

```
RP/0/5/CPU0:R4#Sh bgp vrf CA 5.0.0.0 | inc domain
Sat Apr 21 06:13:34.854 UTC
      Extended community: RT:1:105 OSPF domain-id:0x5:0x010101010101 OSPF
route-type:0:1:0x0 OSPF router-id:4.4.4.4
```

To test and Verify the configuration:

On R1

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
R1#Show ip route ospf | Inc O

R1#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O IA   5.0.0.0/8 [110/66] via 10.1.13.3, 00:07:04, Serial1/0
O E2   6.0.0.0/8 [110/1] via 10.1.13.3, 04:52:47, Serial1/0
O IA      10.1.45.0/24 [110/65] via 10.1.13.3, 00:07:04, Serial1/0
O E2      10.1.46.0/24 [110/1] via 10.1.13.3, 04:52:47, Serial1/0
O IA   192.168.5.0/24 [110/75] via 10.1.13.3, 00:07:04, Serial1/0
O E2   192.168.6.0/24 [110/1] via 10.1.13.3, 00:45:24, Serial1/0
```

Note ONLY the RIP routes are External type 2s and all other OSPF routes are now IA routes.

On R5

```
R5#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O IA   1.0.0.0/8 [110/66] via 10.1.45.4, 00:08:31, GigabitEthernet0/0
O IA      10.1.13.0/24 [110/2] via 10.1.45.4, 00:08:31, GigabitEthernet0/0
O IA   192.168.1.0/24 [110/66] via 10.1.45.4, 00:08:31, GigabitEthernet0/0
```

To verify the configuration;

On R6

```
R6#Sh ip route rip | i R

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route, H - NHRP
R      1.0.0.0/8 [120/5] via 10.1.46.4, 00:00:02, Ethernet0/0
R         10.1.13.0/24 [120/5] via 10.1.46.4, 00:00:02, Ethernet0/0
R      192.168.1.0/24 [120/5] via 10.1.46.4, 00:00:02, Ethernet0/0
```

Since a metric of 5 was set when the routes were redistributed into RIP on R4 (The PE router), all routes end up in R6's routing table with a hop count of 5.

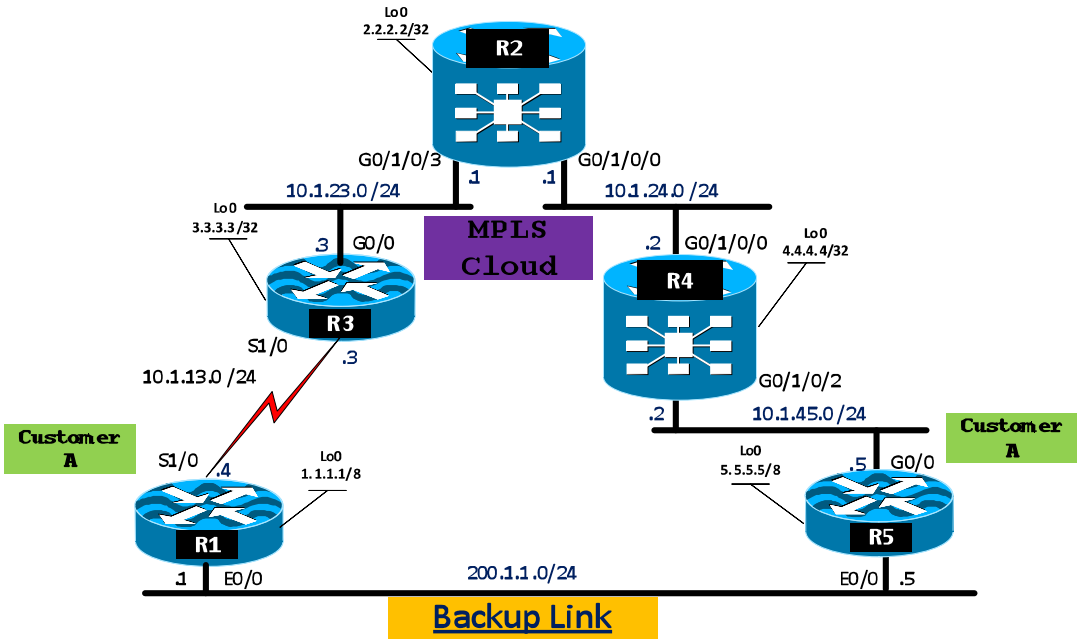
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 4
Backdoor links & OSPF



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 and R5 should be configured in VLAN 100.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the IP addressing according to the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

On R2

```
RP/0/5/CPU0:R2(config)#router ospf 1
RP/0/5/CPU0:R2(config-ospf)#area 0
RP/0/5/CPU0:R2(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R2(config-ospf-ar-if)#passive
RP/0/5/CPU0:R2(config-ospf-ar-if)#commit
RP/0/5/CPU0:R2(config-ospf-ar-if)#end
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 area 0
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 1
RP/0/5/CPU0:R4(config-ospf)#area 0
RP/0/5/CPU0:R4(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R4(config-ospf-ar-if)#passive
RP/0/5/CPU0:R4(config-ospf-ar-if)#commit
RP/0/5/CPU0:R4(config-ospf-ar-if)#end
```

To verify the configuration:

On R3

```
R3#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O        2.2.2.2 [110/2] via 10.1.23.2, 00:01:01, GigabitEthernet0/0
O        4.4.4.4 [110/3] via 10.1.23.2, 00:00:33, GigabitEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O        10.1.24.0/24 [110/2] via 10.1.23.2, 00:01:01, GigabitEthernet0/0
R3#
```

On R2

```
RP/0/5/CPU0:R2#Sh route ospf | i O

O    3.3.3.3/32 [110/2] via 10.1.23.3, 00:01:33, GigabitEthernet0/1/0/3
O    4.4.4.4/32 [110/2] via 10.1.24.4, 00:01:13, GigabitEthernet0/1/0/0
```

On R4

```
RP/0/5/CPU0:R4#Sh route ospf | i O

O    2.2.2.2/32 [110/2] via 10.1.24.2, 00:01:35, GigabitEthernet0/1/0/0
O    3.3.3.3/32 [110/3] via 10.1.24.2, 00:01:35, GigabitEthernet0/1/0/0
O    10.1.23.0/24 [110/2] via 10.1.24.2, 00:01:35, GigabitEthernet0/1/0/0
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 65200 – 65299
R3 – 65300 – 65399
R4 – 65400 – 65499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 65300 65399

R3(config)#int g0/0
R3(config-if)#mpls ip
```

On R2

```
RP/0/5/CPU0:R2(config)#mpls label range 65200 65299
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R2(config)#mpls ldp
RP/0/5/CPU0:R2(config-ldp)#router-id 2.2.2.2
RP/0/5/CPU0:R2(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ldp-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ldp-if)#commit
RP/0/5/CPU0:R2(config-ldp-if)#end
```

On R4

```
RP/0/5/CPU0:R4(config)#mpls label range 65400 65499

RP/0/5/CPU0:R4(config)#mpls ldp
RP/0/5/CPU0:R4(config-ldp)#router-id 4.4.4.4
RP/0/5/CPU0:R4(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ldp-if)#commit
RP/0/5/CPU0:R4(config-ldp-if)#end
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Sh mpls ldp neighbor

Peer LDP Identifier: 2.2.2.2:0
  TCP connection: 2.2.2.2:646 - 4.4.4.4:14570
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 11/10; Downstream-Unsolicited
  Up time: 00:02:12
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    2.2.2.2          10.1.23.2        10.1.24.2
```

On R2

```
RP/0/5/CPU0:R2#Show mpls ldp neighbor

Peer LDP Identifier: 3.3.3.3:0
  TCP connection: 3.3.3.3:57311 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 12/13; Downstream-Unsolicited
  Up time: 00:04:14
  LDP Discovery Sources:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    GigabitEthernet0/1/0/3
  Addresses bound to this peer:
    3.3.3.3          10.1.13.3        10.1.23.3

Peer LDP Identifier: 4.4.4.4:0
  TCP connection: 4.4.4.4:14570 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 11/12; Downstream-Unsolicited
  Up time: 00:02:48
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    4.4.4.4          10.1.24.4        10.1.45.4
```

On R3

```
R3#Sh mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.57311
        State: Oper; Msgs sent/rcvd: 14/13; Downstream
        Up time: 00:05:03
        LDP discovery sources:
          GigabitEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          2.2.2.2         10.1.24.2       10.1.23.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#address-family vpnv4 uni
RP/0/5/CPU0:R4(config-bgp-af)#neighbor 3.3.3.3
RP/0/5/CPU0:R4(config-bgp-nbr)#remote-as 65001
RP/0/5/CPU0:R4(config-bgp-nbr)#update-source lo0
RP/0/5/CPU0:R4(config-bgp-nbr)#address-family vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-nbr-af)#commit
RP/0/5/CPU0:R4(config-bgp-nbr-af)#end
```

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

```
RP/0/5/CPU0:R4#Sh bgp VPNv4 unicast summary | b Neighbor

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
3.3.3.3           0 65001       4       3        1    0    0 00:01:24        0
```

Task 4

Configure a **Virtual Routing Forwarding (VRF) Instance with a name of CA** (For Customer A), a **route-distinguisher (RD) of "1:10"** and a **route-target (RT) of "1:100"** on R3. On R4, the same route targets will be used for the vrf, but the **RD should be configured to be "1:20"** and the name of the VRF should be configured as **CB**.

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-vrf)#route-target Both 1:100

R3(config)#int S1/0
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip address 10.1.13.3 255.255.255.0
```

On R4

```
RP/0/5/CPU0:R4(config)#vrf CB
RP/0/5/CPU0:R4(config-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:100
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:100

RP/0/5/CPU0:R4(config-vrf-af)#int g0/1/0/2
RP/0/5/CPU0:R4(config-if)#No ipv4 address
RP/0/5/CPU0:R4(config-if)#vrf CB
RP/0/5/CPU0:R4(config-if)#ipv4 address 10.1.45.4/24
RP/0/5/CPU0:R4(config-if)#commit
RP/0/5/CPU0:R4(config-if)#end
```

To verify the VRF Configuration:

On R4

```
RP/0/5/CPU0:R4#Sh vrf CB detail

VRF CB; RD not set; VPN ID not set
VRF mode: Regular
Description not set
Interfaces:
  GigabitEthernet0/1/0/2
Address family IPV4 Unicast
  Import VPN route-target communities:
    RT:1:100
  Export VPN route-target communities:
    RT:1:100
  No import route policy
  No export route policy
Address family IPV6 Unicast
  No import VPN route-target communities
  No export VPN route-target communities
  No import route policy
  No export route policy
```

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh ip vrf detail

VRF CA (VRF Id = 1); default RD 1:10; default VPNID <not set>
  Interfaces:
    Se1/0
VRF Table ID = 1
  Export VPN route-target communities
    RT:1:100
  Import VPN route-target communities
    RT:1:100
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Task 5

Configure OSPF using the same process-id on Customer routers (R1 and R5); configure MP-BGP such that Customer A routers can see each others routes. Customer routers should advertise their Lo0 interface with correct mask.

Let's configure the CE routers:

On R1

```
R1(config)#int lo0
R1(config-if)#ip ospf net point-to-point

R1(config)#router ospf 1
R1(config-router)#netw 1.1.1.1 0.0.0.0 area 0
R1(config-router)#netw 10.1.13.1 0.0.0.0 area 0
```

On R5

```
R5(config)#router ospf 1
R5(config-router)#netw 10.1.45.5 0.0.0.0 area 0
R5(config-router)#netw 5.5.5.5 0.0.0.0 area 0

R5(config-router)#int lo0
R5(config-if)#ip ospf net point-to-point
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Let's configure the PE routers:

On R3

```
R3(config)#Router ospf 1 vrf CA
```

OSPF process 1 already exists and is attached to Default-IP-Routing-Table

Note even though the VRF CA is referenced, the same process ID can NOT be used on the same router.

```
R3(config)#router ospf 2 vrf CA
R3(config-router)#Netw 10.1.13.3 0.0.0.0 area 0
```

You should see the following console message:

%OSPF-5-ADJCHG: Process 2, Nbr 1.1.1.1 on Serial0/1 from LOADING to FULL, Loading Done

To verify the configuration:

On R3

```
R3#Sh ip route vrf CA | b Gate

O     1.0.0.0/8 [110/65] via 10.1.13.1, 00:00:03, Serial1/0
      10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C        10.1.13.0/24 is directly connected, Serial1/0
L        10.1.13.3/32 is directly connected, Serial1/0
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 2
RP/0/5/CPU0:R4(config-ospf)#vrf CB
RP/0/5/CPU0:R4(config-ospf-vrf)#area 0
RP/0/5/CPU0:R4(config-ospf-vrf-ar)#int g0/1/0/2
RP/0/5/CPU0:R4(config-ospf-vrf-ar-if)#commit
RP/0/5/CPU0:R4(config-ospf-vrf-ar-if)#end
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Sh route vrf CB | b Gateway

Gateway of last resort is not set
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O    5.0.0.0/8 [110/2] via 10.1.45.5, 00:00:57, GigabitEthernet0/1/0/2
C    10.1.45.0/24 is directly connected, 2d19h, GigabitEthernet0/1/0/2
L    10.1.45.4/32 is directly connected, 2d19h, GigabitEthernet0/1/0/2
```

Redistribute OSPF into MP-BGP on PEs:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#Redistribute OSPF 2
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#rd 1:20
RP/0/5/CPU0:R4(config-bgp-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-bgp-vrf-af)#redistribute ospf 2
```

To verify the configuration:

On R3

```
R3#Sh ip bgp vpnv4 vrf CA | b Netw
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                65         32768 ?
*>i5.0.0.0          4.4.4.4                   2    100      0 ?
*> 10.1.13.0/24     0.0.0.0                   0         32768 ?
*>i10.1.45.0/24     4.4.4.4                   0    100      0 ?
```

On R4

```
RP/0/5/CPU0:R4#Sh ip bgp vrf CB | b Netw

Tue Apr 24 15:37:26.621 UTC
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0/8        3.3.3.3                  65    100      0 ?
*> 5.0.0.0/8        10.1.45.5                 2         32768 ?
*>i10.1.13.0/24     3.3.3.3                   0    100      0 ?
*> 10.1.45.0/24     0.0.0.0                   0         32768 ?
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Processed 4 prefixes, 4 paths
```

Redistribute MP-BGP into OSPF on PEs:

On R3

```
R3(config)#Router ospf 2 vrf CA
R3(config)#domain-id type 0005 vlaue 111111111111
R3(config-router)#Redistribute bgp 65001 subnets
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 2
RP/0/5/CPU0:R4(config-ospf)#vrf CB
RP/0/5/CPU0:R4(config-ospf-vrf)#domain-id type 0005 vlaue 111111111111
RP/0/5/CPU0:R4(config-ospf-vrf)#redistribute bgp 65001
```

To verify the configuration:

On R1

```
R1#sh ip route ospf | i O IA
O IA  5.0.0.0/8 [110/66] via 10.1.13.3, 03:02:51, Serial1/0
O IA     10.1.45.0/24 [110/65] via 10.1.13.3, 03:02:52, Serial1/0

R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```

On R5

```
R5#Sh ip route ospf | i O IA

O IA  1.0.0.0/8 [110/66] via 10.1.45.4, 03:03:32, GigabitEthernet0/0
O IA     10.1.13.0/24 [110/2] via 10.1.45.4, 03:03:32, GigabitEthernet0/0

R5#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/57/60 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Configure the F0/0 interface of the customer routers (R1 and R5) as a backup link between R1 and R5. Ensure that the MPLS Service is preferred over the backup link. Use the existing IP addressing to accomplish this task.

On R1 and R5

```
Rx(config)#int e0/0
Rx(config-if)#NO shut
```

On R1

```
R1(config)#Router ospf 1
R1(config-router)#Netw 200.1.1.1 0.0.0.0 area 0
```

On R5

```
R5(config)#Router ospf 1
R5(config-router)#Netw 200.1.1.5 0.0.0.0 area 0
```

To verify the routing table:

On R5

```
R5#Show ip route ospf | Inc O

O    1.0.0.0/8 [110/2] via 200.1.1.1, 00:00:01, FastEthernet0/0
O       10.1.13.0 [110/65] via 200.1.1.1, 00:00:01, FastEthernet0/0
```

On R1

```
R1#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
O     5.0.0.0/8 [110/11] via 200.1.1.5, 00:00:13, Ethernet0/0
O       10.1.45.0/24 [110/11] via 200.1.1.5, 00:00:13, Ethernet0/0
```

Note the routes are reachable via the backup link. Manipulating OSPF cost will NOT change the routing table, because INTRA-area routes are always preferred over INTER-area routes. One way to fix this problem is to configure a Sham link; Sham link was created to resolve this problem. Sham link is an OSPF INTRA-area link configured between the two PE routers. Sham link is included in the SPF calculation, just like any link in OSPF.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

When LSAs are flooded across this link all OSPF route types are preserved and they don't have to be converted into INTRA-area or External.

To configure a Sham link:

5. **Create a loopback interface and assign an IP address with a prefix-length of 32.**

6. **Enable VRF forwarding on the loopback interface, in this case "IP vrf forwarding CA" on R3 and "IP vrf forwarding CB" on R4.**

7. **Advertise the /32 IP address of the loopback interface in BGP, on R3 it should be advertised under "Address-family ipv4 vrf CA" and on R4 it should be advertised under "Address-family ipv4 vrf CB".**

8. **Configure Sham link under OSPF, on R3 it should be done under "router ospf 2 vrf CA" and on R4 it should be done under "router ospf 2 vrf CB".**

Step One:

Configure the loopback with a prefix-length of 32:

On R3

```
R3(config-router)#int lo34
R3(config-if)#ip addr 34.1.1.3 255.255.255.255
```

On R4

```
RP/0/5/CPU0:R4(config)#int lo43
RP/0/5/CPU0:R4(config-if)#ipv4 address 34.1.1.4/32
RP/0/5/CPU0:R4(config-if)#commit
RP/0/5/CPU0:R4(config-if)#root
```

Step Two:

Configure VRF forwarding on the loopback interface:

On R3

```
R3(config)#int lo34
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip addr 34.1.1.3 255.255.255.255
```

On R4

```
RP/0/5/CPU0:R4(config)#int lo43
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config-if)#NO ipv4 address
RP/0/5/CPU0:R4(config-if)#vrf CB
RP/0/5/CPU0:R4(config-if)#ipv4 address 34.1.1.4/32
RP/0/5/CPU0:R4(config-if)#commit
RP/0/5/CPU0:R4(config-if)#root
```

Step Three:

Advertise the ip address of the loopback interface in MP-BGP:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#network 34.1.1.3 mask 255.255.255.255
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#address-family ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-af)#netw 34.1.1.4/32
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
RP/0/5/CPU0:R4(config-bgp-vrf-af)#root
```

Step Four:

Configure the Sham link:

On R3

```
R3(config)#router ospf 2 vrf CA
R3(config-router)#Area 0 Sham-link 34.1.1.3 34.1.1.4 cost 1
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 2
RP/0/5/CPU0:R4(config-ospf)#vrf CB
RP/0/5/CPU0:R4(config-ospf-vrf)#area 0 Sham-link 34.1.1.4 34.1.1.3 cost 1
RP/0/5/CPU0:R4(config-ospf-vrf)#commit
RP/0/5/CPU0:R4(config-ospf-vrf)#end
```

You should see the following console message on R3:

***Mar 4 06:16:40.556: %OSPF-5-ADJCHG: Process 2, Nbr 4.4.4.4 on OSPF_SL1 from LOADING to FULL, Loading Done**

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

To verify the configuration:

On R3

```
R3#Sh ip ospf sham-links

Sham Link OSPF_SL1 to address 34.1.1.4 is up
Area 0 source address 34.1.1.3
  Run as demand circuit
  DoNotAge LSA allowed. Cost of using 1 State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
    Hello due in 00:00:02
    Adjacency State FULL (Hello suppressed)
    Index 2/2, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

Area 0 was the source for the sham link. This will commonly match the area that has been connected from the PE to the CE. This connection is considered "up" and operational at this time.

As you can see, this connection will run as a demand circuit, The route will remain up indefinately with no interference. The "DoNotAge LSA allowed" should back up this statement with type point-to-point and a cost of 1 as configured under the OSPF process for the VPN. The remaining timers reflect normal OSPF peering.

On R3

```
R3#Sh ip ospf neighbor


Neighbor ID     Pri   State           Dead Time   Address         Interface
2.2.2.2           1   FULL/DR         00:00:39    10.1.23.2
GigabitEthernet0/0
4.4.4.4           0   FULL/  -            -        34.1.1.4        OSPF_SL1
1.1.1.1           0   FULL/  -        00:00:32    10.1.13.1       Serial1/0
```

On R1

```
R1#Sh ip route ospf | i O

        D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
        N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
        E1 - OSPF external type 1, E2 - OSPF external type 2
        o - ODR, P - periodic downloaded static route, H - NHRP
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O     5.0.0.0/8 [110/11] via 200.1.1.5, 01:19:12, Ethernet0/0
O        10.1.45.0/24 [110/11] via 200.1.1.5, 01:19:12, Ethernet0/0
O E2     34.1.1.3 [110/1] via 200.1.1.5, 01:04:24, Ethernet0/0
O E2     34.1.1.4 [110/1] via 200.1.1.5, 00:04:37, Ethernet0/0
```

On R4

```
RP/0/5/CPU0:R4#Sh ospf vrf CB sham-links

Sham Links for OSPF 2, VRF CB

Sham Link OSPF_SL0 to address 34.1.1.4 is down
Area 0, source address 34.1.1.4
IfIndex = 2
  Run as demand circuit
  DoNotAge LSA allowed., Cost of using 1
  Transmit Delay is 1 sec, State DOWN,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Sham Link OSPF_SL1 to address 34.1.1.3 is up
Area 0, source address 34.1.1.4
IfIndex = 2
  Run as demand circuit
  DoNotAge LSA allowed., Cost of using 1
  Transmit Delay is 1 sec, State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:06
    Adjacency State FULL (Hello suppressed)
    Number of DBD retrans during last exchange 0
    Index 2/2, retransmission queue length 0, number of retransmission 0
    First 0(0)/0(0) Next 0(0)/0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec

RP/0/5/CPU0:R4#Sh ospf vrf CB neighbor | b Neighbor

Neighbors for OSPF 2, VRF CB

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.13.3       1     FULL/  -        00:00:36    34.1.1.3        OSPF_SL1
    Neighbor is up for 00:00:03
5.5.5.5         1     FULL/DR         00:00:38    10.1.45.5
GigabitEthernet0/1/0/2
    Neighbor is up for 07:42:36

Total neighbor count: 2
```

On R1

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R1#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O     5.0.0.0/8 [110/11] via 200.1.1.5, 00:00:18, Ethernet0/0
O        10.1.45.0/24 [110/11] via 200.1.1.5, 00:00:18, Ethernet0/0
O E2     34.1.1.3 [110/1] via 200.1.1.5, 00:00:18, Ethernet0/0
O E2     34.1.1.4 [110/1] via 200.1.1.5, 00:00:18, Ethernet0/0
```

Note the reason the routers still prefer the backup link is because of the OSPF cost, note the routes are no longer an INTER-area route and they have preserved their true route type.
Since the router sees the routes through the provider and the backup link as INTRA-area routes, the cost can be manipulated such that the routers will take the provider as their primary and the backup link will be used ONLY if the primary link is down.

In order to see the actual cost through the provider, the backup link should be shutdown:

On R1 and R5

```
Rx(config)#Int e0/0
Rx(config-if)#Shut
```

The following show command reveals the OSPF cost through the provider:

On R5

```
R5#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O     1.0.0.0/8 [110/67] via 10.1.45.4, 00:01:30, GigabitEthernet0/0
O        10.1.13.0/24 [110/66] via 10.1.45.4, 00:01:30, GigabitEthernet0/0
O E2     34.1.1.3 [110/1] via 10.1.45.4, 00:01:30, GigabitEthernet0/0
O E2     34.1.1.4 [110/1] via 10.1.45.4, 00:01:30, GigabitEthernet0/0
```

In order to force the provider link as the primary and the FastEthernet link as the backup link, a cost of 131 is assigned to the backup link, as follows:

On Both Routers

```
Rx(config)#Int F0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Rx(config-if)#ip ospf cost 131
Rx(config-if)#NO Shut
```

To verify the configuration:

On R1

```
R1#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O     5.0.0.0/8 [110/67] via 10.1.13.3, 00:02:12, Serial1/0
O       10.1.45.0/24 [110/66] via 10.1.13.3, 00:02:12, Serial1/0
O E2    34.1.1.3 [110/1] via 10.1.13.3, 00:02:12, Serial1/0
O E2    34.1.1.4 [110/1] via 10.1.13.3, 00:02:12, Serial1/0
```

On R5

```
R5#Sh ip route ospf | i O

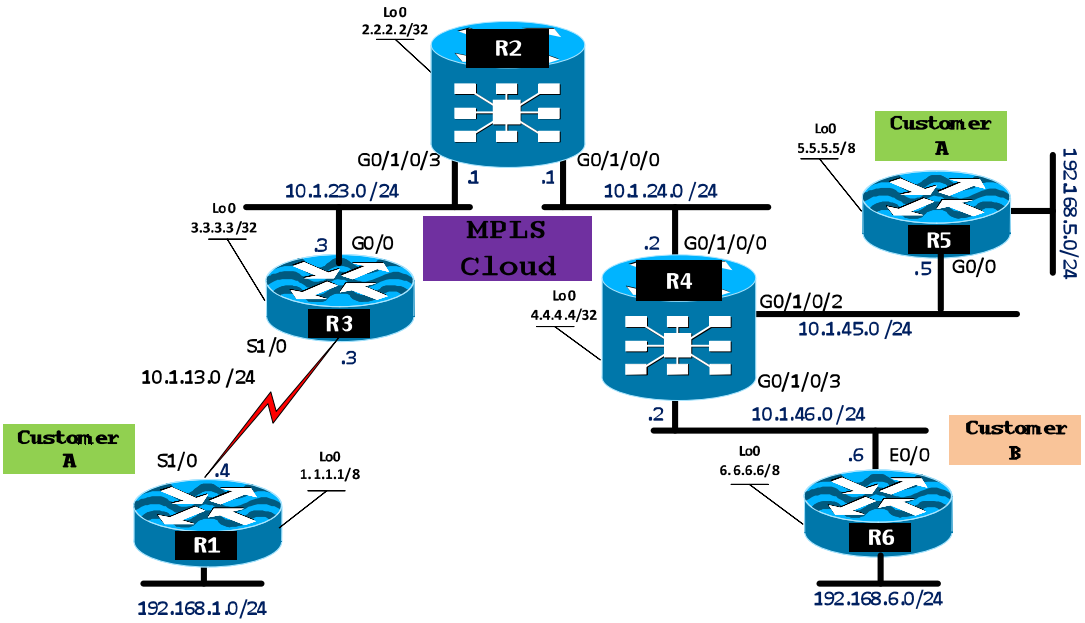
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       o - ODR, P - periodic downloaded static route, H - NHRP
O     1.0.0.0/8 [110/67] via 10.1.45.4, 00:03:17, GigabitEthernet0/0
O       10.1.13.0/24 [110/66] via 10.1.45.4, 00:03:17, GigabitEthernet0/0
O E2    34.1.1.3 [110/1] via 10.1.45.4, 00:03:17, GigabitEthernet0/0
O E2    34.1.1.4 [110/1] via 10.1.45.4, 00:03:17, GigabitEthernet0/0
```

Task 7

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 5
Eigrp Routing in a VPN



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100, R5 in VLAN 500 and F0/1 interface of R6 in VLAN 600.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the F0/1 interface of R6 in VLAN 6.
- Configure the F0/0 interface of R4 and R6 in VLAN 900
- Configure the rest of the routers according to the above diagram.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

On R2

```
RP/0/5/CPU0:R2(config)#router ospf 1
RP/0/5/CPU0:R2(config-ospf)#area 0
RP/0/5/CPU0:R2(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R2(config-ospf-ar-if)#passive
RP/0/5/CPU0:R2(config-ospf-ar-if)#commit
RP/0/5/CPU0:R2(config-ospf-ar-if)#root
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 1
RP/0/5/CPU0:R4(config-ospf)#area 0
RP/0/5/CPU0:R4(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R4(config-ospf-ar-if)#passive
RP/0/5/CPU0:R4(config-ospf-ar-if)#commit
RP/0/5/CPU0:R4(config-ospf-ar-if)#root
```

To verify the configuration:

On R3

```
R3#Sh ip route ospf | i O_

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
O        2.2.2.2 [110/2] via 10.1.23.2, 00:01:46, GigabitEthernet0/0
O        4.4.4.4 [110/3] via 10.1.23.2, 00:00:50, GigabitEthernet0/0
O        10.1.24.0/24 [110/2] via 10.1.23.2, 00:01:46, GigabitEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
RP/0/5/CPU0:R2#Sh route ospf | i O

O    3.3.3.3/32 [110/2] via 10.1.23.3, 00:02:56, GigabitEthernet0/1/0/3
O    4.4.4.4/32 [110/2] via 10.1.24.4, 00:01:58, GigabitEthernet0/1/0/0
```

On R4

```
RP/0/5/CPU0:R4#Sh route ospf | i O

O    2.2.2.2/32 [110/2] via 10.1.24.2, 00:02:32, GigabitEthernet0/1/0/0
O    3.3.3.3/32 [110/3] via 10.1.24.2, 00:02:32, GigabitEthernet0/1/0/0
O    10.1.23.0/24 [110/2] via 10.1.24.2, 00:02:32, GigabitEthernet0/1/0/0
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 65200 – 65299
R3 – 65300 – 65399
R4 – 65400 – 65499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 300 399

R3(config)#int g0/0
R3(config-if)#MPLS IP
```

On R2

```
RP/0/5/CPU0:R2(config)#mpls label range 65200 65299
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R2(config)#mpls ldp
RP/0/5/CPU0:R2(config-ldp)#router-id 2.2.2.2

RP/0/5/CPU0:R2(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ldp-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ldp-if)#commit
RP/0/5/CPU0:R2(config-ldp-if)#end
```

On R4

```
RP/0/5/CPU0:R4(config)#mpls label range 65400 65499

RP/0/5/CPU0:R4(config)#mpls ldp
RP/0/5/CPU0:R4(config-ldp)#router-id 4.4.4.4

R4(config)#MPLS label range 400 499

RP/0/5/CPU0:R4(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ldp-if)#commit
RP/0/5/CPU0:R4(config-ldp-if)#end
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Sh mpls ldp neighbor

Peer LDP Identifier: 2.2.2.2:0
  TCP connection: 2.2.2.2:646 - 4.4.4.4:25673
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 12/10; Downstream-Unsolicited
  Up time: 00:01:59
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    2.2.2.2          10.1.23.2        10.1.24.2
```

On R2

```
RP/0/5/CPU0:R2#Sh mpls ldp neighbor

Peer LDP Identifier: 3.3.3.3:0
  TCP connection: 3.3.3.3:13503 - 2.2.2.2:646
  Graceful Restart: No
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 15/16; Downstream-Unsolicited
  Up time: 00:06:27
  LDP Discovery Sources:
    GigabitEthernet0/1/0/3
  Addresses bound to this peer:
    3.3.3.3          10.1.13.3        10.1.23.3

Peer LDP Identifier: 4.4.4.4:0
  TCP connection: 4.4.4.4:25673 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 11/13; Downstream-Unsolicited
  Up time: 00:03:07
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    4.4.4.4          10.1.24.4        10.1.45.4        10.1.46.4
```

On R3

```
R3#Sh mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.13503
        State: Oper; Msgs sent/rcvd: 17/16; Downstream
        Up time: 00:07:24
        LDP discovery sources:
          GigabitEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          2.2.2.2         10.1.24.2       10.1.23.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#address vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-af)#neighbor 3.3.3.3
RP/0/5/CPU0:R4(config-bgp-nbr)#remote-as 65001
RP/0/5/CPU0:R4(config-bgp-nbr)#update-source lo0
RP/0/5/CPU0:R4(config-bgp-nbr)#address-family vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-nbr-af)#commit
RP/0/5/CPU0:R4(config-bgp-nbr-af)#end
```

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

```
RP/0/5/CPU0:R4#sh bgp VPNv4 unicast summary | b Neighbor
Thu Apr 26 03:44:43.914 UTC
Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
3.3.3.3           0 65001       4       3        1    0    0 00:01:35           0
```

Task 4

Configure VRFs on R3 and R4 and enable VRF forwarding on the interfaces of these two routers based on the following chart.

Router	VRF	RD	RT	Interface
R3	CA	1:10	1:105	S1/0
R4	CA	1:50	1:105	G0/1/0/2
	CB	1:60	1:600	G0/1/0/3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#rd 1:10
R3(config-vrf)#route-target both 1:105

R3(config-vrf)#int S1/0
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip addr 10.1.13.3 255.255.255.0
```

On R4

```
RP/0/5/CPU0:R4(config)#vrf CA
RP/0/5/CPU0:R4(config-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:105
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:105
RP/0/5/CPU0:R4(config-vrf-af)#commit
RP/0/5/CPU0:R4(config-vrf-af)#root

RP/0/5/CPU0:R4(config)#vrf CB
RP/0/5/CPU0:R4(config-vrf)#address-family ipv4 unicast
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:600
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:600
RP/0/5/CPU0:R4(config-vrf-af)#commit
RP/0/5/CPU0:R4(config-vrf-af)#root

RP/0/5/CPU0:R4(config)#int g0/1/0/2
RP/0/5/CPU0:R4(config-if)#no ipv4 address
RP/0/5/CPU0:R4(config-if)#vrf CA
RP/0/5/CPU0:R4(config-if)#ipv4 address 10.1.45.4/24
RP/0/5/CPU0:R4(config-if)#commit

RP/0/5/CPU0:R4(config-if)#int g0/1/0/3
RP/0/5/CPU0:R4(config-if)#no ipv4 address
RP/0/5/CPU0:R4(config-if)#vrf CB
RP/0/5/CPU0:R4(config-if)#ipv4 address 10.1.46.4/24
RP/0/5/CPU0:R4(config-if)#commit
RP/0/5/CPU0:R4(config-if)#end
```

To verify the configuration:

On R3

```
R3#Sh ip vrf det
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
VRF CA (VRF Id = 1); default RD 1:10; default VPNID <not set>
  Interfaces:
    Se1/0
VRF Table ID = 1
  Export VPN route-target communities
    RT:1:105
  Import VPN route-target communities
    RT:1:105
  No import route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

On R4

```
RP/0/5/CPU0:R4#Sh vrf CA detail


VRF CA; RD not set; VPN ID not set
VRF mode: Regular
Description not set
Interfaces:
  GigabitEthernet0/1/0/2
Address family IPV4 Unicast
  Import VPN route-target communities:
    RT:1:105
  Export VPN route-target communities:
    RT:1:105
  No import route policy
  No export route policy
Address family IPV6 Unicast
  No import VPN route-target communities
  No export VPN route-target communities
  No import route policy
  No export route policy
```

Remember, the RD is a BGP parameter and is set when the vrf is placed under the BGP process in IOS-XR.

```
RP/0/5/CPU0:R4#Sh vrf CB detail

VRF CB; RD not set; VPN ID not set
VRF mode: Regular
Description not set
Interfaces:
  GigabitEthernet0/1/0/3
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Address family IPV4 Unicast
  Import VPN route-target communities:
    RT:1:600
  Export VPN route-target communities:
    RT:1:600
  No import route policy
  No export route policy
Address family IPV6 Unicast
  No import VPN route-target communities
  No export VPN route-target communities
  No import route policy
  No export route policy  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
```

Task 5

Configure Customer A routers (R1 and R5) with a VRF service that incorporates Eigrp as the routing protocol; Configure the VRF applied to R6 to use OSPF and ensure that this customer participates in the same Closed User Group (CUG) as R1 and R5.
R6 should advertise its loopback0 interface with the correct mask.

To configure the CE Routers:

On R1

```
R1(config)#router eigrp 100
R1(config-router)#No au
R1(config-router)#network 10.1.13.1 0.0.0.0
R1(config-router)#netw 1.1.1.1 0.0.0.0
```

On R5

```
R5(config)#router eigrp 100
R5(config-router)#No au
R5(config-router)#netw 10.1.45.5 0.0.0.0
R5(config-router)#netw 5.5.5.5 0.0.0.0
```

On R6

```
R6(config-router)#int lo0
R6(config-if)#ip ospf net point-to-point
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R6(config)#router ospf 1
R6(config-router)#netw 10.1.46.6 0.0.0.0 are 0
R6(config-router)#netw 6.6.6.6 0.0.0.0 are 0
```

Configuring the PE Routers:

On R3

Configuring Eigrp between PE-CE routers varies from the traditional Eigrp configuration, notice that the definition of Autonomous system number (AS#) is configured under the address-family for the VRF, in this case CA, the AS number configured under the address-family MUST match the AS number configured on the CE router/s. Just like RIPv2 configuration, multiple Eigrp routing context can be configured for every VRF using the address-family command.
One point to remember is that the passive-interface cannot be used in the global instance of EIGRP because it will be inherited to the VPN when the VRF is enabled on the interface and will adversely affect route propagation.

```
R3(config)#Router eigrp 100
R3(config-router)#Address-family ipv4 vrf CA
R3(config-router-af)#No au
R3(config-router-af)#autonomous-system 100
R3(config-router-af)#network 10.1.13.3 0.0.0.0
```



You should get the following console message after adding the "Autonomous-system" command:

%DUAL-5-NBRCHANGE: IP-EIGRP(1) 100: Neighbor 10.1.13.1 (Serial1/0) is up: new adjacency

To verify the configuration:

On R3

```
R3#Sh ip route vrf CA | b Gateway
Gateway of last resort is not set

D      1.0.0.0/8 [90/2297856] via 10.1.13.1, 00:00:54, Serial1/0
       10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C         10.1.13.0/24 is directly connected, Serial1/0
L         10.1.13.3/32 is directly connected, Serial1/0
```

On R4

```
RP/0/5/CPU0:R4(config)#router eigrp 100
RP/0/5/CPU0:R4(config-eigrp)#vrf CA
RP/0/5/CPU0:R4(config-eigrp-vrf)#address-family ipv4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#autonomous-system 100
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#interface GigabitEthernet0/1/0/2
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#commit
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#end
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Sh route vrf CA | b Gateway

Gateway of last resort is not set

D    5.0.0.0/8 [90/130816] via 10.1.45.5, 00:02:46,
GigabitEthernet0/1/0/2
C    10.1.45.0/24 is directly connected, 00:29:28, GigabitEthernet0/1/0/2
L    10.1.45.4/32 is directly connected, 00:29:28, GigabitEthernet0/1/0/2
```

Configuring the PE router (R4) for Customer B (R6):

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 2
RP/0/5/CPU0:R4(config-ospf)#vrf CB
RP/0/5/CPU0:R4(config-ospf-vrf)#area 0
RP/0/5/CPU0:R4(config-ospf-vrf-ar)#interface GigabitEthernet0/1/0/3
RP/0/5/CPU0:R4(config-ospf-vrf-ar-if)#commit
RP/0/5/CPU0:R4(config-ospf-vrf-ar-if)#end
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Show route vrf CB | b Gateway

Gateway of last resort is not set

O    6.0.0.0/8 [110/2] via 10.1.46.6, 00:02:03, GigabitEthernet0/1/0/3
C    10.1.46.0/24 is directly connected, 00:33:03, GigabitEthernet0/1/0/3
L    10.1.46.4/32 is directly connected, 00:33:03, GigabitEthernet0/1/0/3
```

Perform a mutual redistribution on the PE and CE routers and verify the routes:

On R3

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config)#Router bgp 65001
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#redistribute Eigrp 100

R3(config)#Router eigrp 100
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#Redistribute BGP 65001 metric 1 1 1 1 1
```

To verify the configuration:

On R3

```
R3#Sh ip bgp vpnv4 vrf CA | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1          2297856         32768 ?
*> 10.1.13.0/24     0.0.0.0                  0         32768 ?
```

On R4

```
RP/0/5/CPU0:R4(config)#router eigrp 100
RP/0/5/CPU0:R4(config-eigrp)#vrf CA
RP/0/5/CPU0:R4(config-eigrp-vrf)#address ipv4
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#redistribute bgp 65001
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#root

RP/0/5/CPU0:R4(config)#router ospf 2
RP/0/5/CPU0:R4(config-ospf)#vrf CB
RP/0/5/CPU0:R4(config-ospf-vrf)#redistribute bgp 65001
RP/0/5/CPU0:R4(config-ospf-vrf)#commit
RP/0/5/CPU0:R4(config-ospf-vrf)#root

RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CA
RP/0/5/CPU0:R4(config-bgp-vrf)#rd 1:10
RP/0/5/CPU0:R4(config-bgp-vrf)#address ipv4 unicast
RP/0/5/CPU0:R4(config-bgp-vrf-af)#redistribute eigrp 100
RP/0/5/CPU0:R4(config-bgp-vrf-af)#exit
RP/0/5/CPU0:R4(config-bgp-vrf)#exit

RP/0/5/CPU0:R4(config-bgp)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#rd 1:20
RP/0/5/CPU0:R4(config-bgp-vrf)#address ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-af)#redistribute ospf 2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
RP/0/5/CPU0:R4(config-bgp-vrf-af)#end
```

To verify the configuration:

On R5

```
R5#Sh ip route eigrp | i D

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route, H - NHRP
D     1.0.0.0/8 [90/2298112] via 10.1.45.4, 00:01:23, GigabitEthernet0/0
D       10.1.13.0/24 [90/2170112] via 10.1.45.4, 00:01:23, GigabitEthernet0/0
```

On R1

```
R1#Sh ip route eigrp | inc D

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route, H - NHRP
D     5.0.0.0/8 [90/2298112] via 10.1.13.3, 00:02:54, Serial1/0
D       10.1.45.0/24 [90/2170112] via 10.1.13.3, 00:02:54, Serial1/0
```

On R3

```
R3#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1          2297856         32768 ?
*>i5.0.0.0          4.4.4.4             130816    100      0 ?
*> 10.1.13.0/24     0.0.0.0                  0         32768 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
Route Distinguisher: 1:50
*>i5.0.0.0          4.4.4.4             130816    100      0 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
```

On R4

```
RP/0/5/CPU0:R4#Sh bgp vpnv4 unicast | b Netw

   Network            Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10
*>i1.0.0.0/8         3.3.3.3            2297856    100      0 ?
*>i10.1.13.0/24      3.3.3.3                  0    100      0 ?
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Route Distinguisher: 1:50 (default for vrf CA)
*>i1.0.0.0/8          3.3.3.3             2297856    100      0 ?
*> 5.0.0.0/8          10.1.45.5            130816         32768 ?
*>i10.1.13.0/24       3.3.3.3                   0    100      0 ?
*> 10.1.45.0/24       0.0.0.0                   0         32768 ?
Route Distinguisher: 1:60 (default for vrf CB)
*> 6.0.0.0/8          10.1.46.6                 2         32768 ?
*> 10.1.46.0/24       0.0.0.0                   0         32768 ?
```

Note R3 contains networks that are associated to VRF CA ONLY, and it does NOT contain any information about VRF CB. Is there something wrong?

When R4 sends all of its routes and their associated route-targets to its PE peer (R3), R3 checks its VRF configurations, and if it does NOT see a given route-target configuration locally, it will simply filter them. It's very similar to VTP pruning. To avoid this behavior, R3 can be configured NOT to filter these routes:

On R3

```
R3(config)#router bgp 65001
R3(config-router)#NO bgp default route-target filter

R3#Clear ip bgp * in
```

To verify the configuration:

On R3

```
R3#Sh ip bgp vpnv4 all | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1          2297856         32768 ?
*>i5.0.0.0          4.4.4.4             130816    100      0 ?
*> 10.1.13.0/24     0.0.0.0                  0         32768 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
Route Distinguisher: 1:50
*>i5.0.0.0          4.4.4.4             130816    100      0 ?
*>i10.1.45.0/24     4.4.4.4                  0    100      0 ?
Route Distinguisher: 1:60
*>i6.0.0.0          4.4.4.4                  2    100      0 ?
*>i10.1.46.0/24     4.4.4.4                  0    100      0 ?
```

To test the configuration:

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R5

```
R5#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

The last configuration in this task requires that R6 to participate in the same Closed User Group (CUG) as R1 and R5:

On R4

```
RP/0/5/CPU0:R4(config)#vrf CB
RP/0/5/CPU0:R4(config-vrf)#address ipv4 uni
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:105
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:105
RP/0/5/CPU0:R4(config-vrf-af)#commit
```

Because the router type is different than EIGRP, Te IOS-XR router finally warns us that a default metric is required. If a metric is not selected, R5 will not see the routes that are generated by R6.

```
: eigrp[1002]: IP-EIGRP(100)-v4: ipigrp2_pick_redist_metric: No Default
metrics set

RP/0/5/CPU0:R4(config)#router eigrp 100
RP/0/5/CPU0:R4(config-eigrp)#vrf CA
RP/0/5/CPU0:R4(config-eigrp-vrf)#address ipv4
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#default-metric 1 1 1 1 1
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#commit
RP/0/5/CPU0:R4(config-eigrp-vrf-af)#end
```

To verify test the configuration:

On R6

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R6#Sh ip route ospf | i O_

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
O E2  1.0.0.0/8 [110/1] via 10.1.46.4, 00:01:35, Ethernet0/0
O E2  5.0.0.0/8 [110/1] via 10.1.46.4, 00:01:35, Ethernet0/0
O E2     10.1.13.0/24 [110/1] via 10.1.46.4, 00:01:35, Ethernet0/0
O E2     10.1.45.0/24 [110/1] via 10.1.46.4, 00:01:35, Ethernet0/0

R6#Ping 1.1.1.1

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms

R6#Ping 5.5.5.5

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 5.5.5.5, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On R1

```
R1#Sh ip route eigrp | i D

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       o - ODR, P - periodic downloaded static route, H - NHRP
D     5.0.0.0/8 [90/2298112] via 10.1.13.3, 00:21:13, Serial1/0
D EX  6.0.0.0/8 [170/2560512256] via 10.1.13.3, 00:09:45, Serial1/0
D        10.1.45.0/24 [90/2170112] via 10.1.13.3, 00:21:13, Serial1/0
D EX     10.1.46.0/24 [170/2560512256] via 10.1.13.3, 00:09:45, Serial1/0

R1#Ping 6.6.6.6

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

On R5

```
R5#Sh ip route eigrp | i D_

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
D     1.0.0.0/8 [90/2298112] via 10.1.45.4, 00:22:46, GigabitEthernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
D EX  6.0.0.0/8 [170/2560000512] via 10.1.45.4, 00:07:40, GigabitEthernet0/0
D        10.1.13.0/24 [90/2170112] via 10.1.45.4, 00:22:46, GigabitEthernet0/0
D EX     10.1.46.0/24

R5#Ping 6.6.6.6

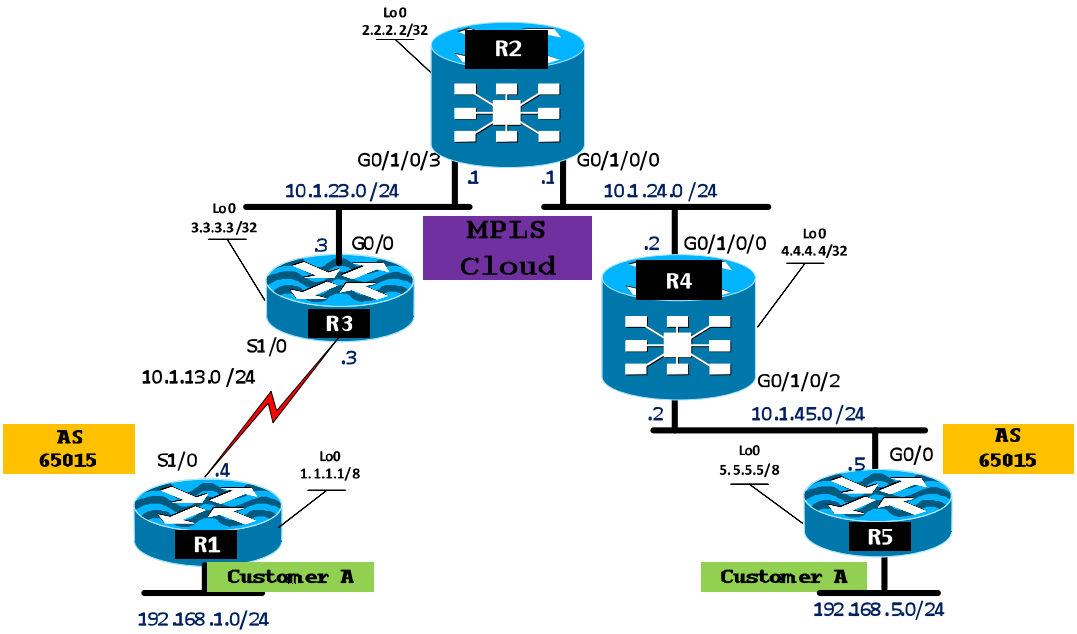
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 6.6.6.6, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 6
BGP Routing in a VPN



Lab Setup:

- The connections between R1 and R3, R4 and R5 should be configured with HDLC encapsulation. The clock rate should be set to 64000.
- Configure F0/0 of R1 in VLAN 100 and R5 in VLAN 500.
- Configure the F0/0 interface of R2 and R3 in VLAN 200.
- Configure the F0/1 interface of R2 and R4 in VLAN 300.
- Configure the rest of the routers according to the above diagram.

Task 1

Configure OSPF on Core MPLS routers (R2, R3 and R4), you should run OSPF area 0 on the F0/0 interface of R2 and R3, F0/1 interface of R2 and R4 and the loopback interfaces of these routers.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R2

```
RP/0/5/CPU0:R2(config)#router ospf 1
RP/0/5/CPU0:R2(config-ospf)#area 0
RP/0/5/CPU0:R2(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R2(config-ospf-ar-if)#passive
RP/0/5/CPU0:R2(config-ospf-ar-if)#commit
RP/0/5/CPU0:R2(config-ospf-ar-if)#end
```

On R3

```
R3(config)#router ospf 1
R3(config-router)#netw 3.3.3.3 0.0.0.0 area 0
R3(config-router)#netw 10.1.23.3 0.0.0.0 are 0
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 1
RP/0/5/CPU0:R4(config-ospf)#area 0
RP/0/5/CPU0:R4(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R4(config-ospf-ar-if)#passive
RP/0/5/CPU0:R4(config-ospf-ar-if)#commit
RP/0/5/CPU0:R4(config-ospf-ar-if)#end
```

To verify the configuration:

On R3

```
R3#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
O        2.2.2.2 [110/2] via 10.1.23.2, 00:03:14, GigabitEthernet0/0
O        4.4.4.4 [110/3] via 10.1.23.2, 00:02:22, GigabitEthernet0/0
O        10.1.24.0/24 [110/2] via 10.1.23.2, 00:03:14, GigabitEthernet0/0
```

On R2

```
RP/0/5/CPU0:R2#Sh route ospf | i O

O    3.3.3.3/32 [110/2] via 10.1.23.3, 00:03:59, GigabitEthernet0/1/0/3
O    4.4.4.4/32 [110/2] via 10.1.24.4, 00:03:11, GigabitEthernet0/1/0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R4

```
RP/0/5/CPU0:R4#Sh route ospf | i O

O    2.2.2.2/32 [110/2] via 10.1.24.2, 00:03:47, GigabitEthernet0/1/0/0
O    3.3.3.3/32 [110/3] via 10.1.24.2, 00:03:47, GigabitEthernet0/1/0/0
O    10.1.23.0/24 [110/2] via 10.1.24.2, 00:03:47, GigabitEthernet0/1/0/0
```

Task 2

Configure LDP between the core routers. These routers should use their loopback 0 interfaces as their LDP router ID; the core MPLS routers (R2, R3 and R4) should use the following label range:

R2 – 65200 – 65299
R3 – 65300 – 65399
R4 – 65400 – 65499

On R3

```
R3(config)#MPLS label protocol LDP

R3(config)#MPLS ldp router-id lo0

R3(config)#MPLS label range 65300 65399

R3(config)#int g0/0
R3(config-if)#MPLS IP
```

On R2

```
RP/0/5/CPU0:R2(config)#mpls label range 65200 65299

RP/0/5/CPU0:R2(config)#mpls ldp
RP/0/5/CPU0:R2(config-ldp)#router-id 2.2.2.2
RP/0/5/CPU0:R2(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R2(config-ldp-if)#int g0/1/0/3
RP/0/5/CPU0:R2(config-ldp-if)#commit
RP/0/5/CPU0:R2(config-ldp-if)#end
```

On R4

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config)#mpls label range 65400 65499

RP/0/5/CPU0:R4(config)#mpls ldp
RP/0/5/CPU0:R4(config-ldp)#router-id 4.4.4.4
RP/0/5/CPU0:R4(config-ldp)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ldp-if)#commit
RP/0/5/CPU0:R4(config-ldp-if)#end
```

To verify the configuration:

On R4

```
RP/0/5/CPU0:R4#Sh mpls ldp neighbor

Peer LDP Identifier: 2.2.2.2:0
  TCP connection: 2.2.2.2:646 - 4.4.4.4:32995
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 10/9; Downstream-Unsolicited
  Up time: 00:01:14
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    2.2.2.2          10.1.23.2        10.1.24.2
```

On R2

```
RP/0/5/CPU0:R2#Sh mpls ldp neighbor

Peer LDP Identifier: 3.3.3.3:0
  TCP connection: 3.3.3.3:16107 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 10/11; Downstream-Unsolicited
  Up time: 00:02:05
  LDP Discovery Sources:
    GigabitEthernet0/1/0/3
  Addresses bound to this peer:
    3.3.3.3          10.1.13.3        10.1.23.3

Peer LDP Identifier: 4.4.4.4:0
  TCP connection: 4.4.4.4:32995 - 2.2.2.2:646
  Graceful Restart: No
  Session Holdtime: 180 sec
  State: Oper; Msgs sent/rcvd: 10/11; Downstream-Unsolicited
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Up time: 00:02:01
  LDP Discovery Sources:
    GigabitEthernet0/1/0/0
  Addresses bound to this peer:
    4.4.4.4          10.1.24.4        10.1.45.4
```

On R3

```
R3#Sh mpls ldp neighbor

    Peer LDP Ident: 2.2.2.2:0; Local LDP Ident 3.3.3.3:0
        TCP connection: 2.2.2.2.646 - 3.3.3.3.16107
        State: Oper; Msgs sent/rcvd: 12/11; Downstream
        Up time: 00:02:58
        LDP discovery sources:
          GigabitEthernet0/0, Src IP addr: 10.1.23.2
        Addresses bound to peer LDP Ident:
          2.2.2.2         10.1.24.2       10.1.23.2
```

Task 3

Configure MP-BGP between R3 and R4 as they represent the Provider Edge routers in this topology in AS 65001. Do not allow the BGP peers to share IPV4 routing information by default. The only bgp peering relationship should be VPNv4.

On R3

```
R3(config)#Router bgp 65001
R3(config-router)#NO BGP default ipv4-unicast
R3(config-router)#Neighbor 4.4.4.4 remote-as 65001
R3(config-router)#Neighbor 4.4.4.4 Update-source Lo0

R3(config-router)#Address-family vpnv4 unicast
R3(config-router-af)#Neighbor 4.4.4.4 Activate
R3(config-router-af)#Neighbor 4.4.4.4 Send-community extended
```

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#address-family vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-af)#neighbor 3.3.3.3
RP/0/5/CPU0:R4(config-bgp-nbr)#remote-as 65001
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config-bgp-nbr)#update-source lo0
RP/0/5/CPU0:R4(config-bgp-nbr)#address-family vpnv4 unicast
RP/0/5/CPU0:R4(config-bgp-nbr-af)#commit
RP/0/5/CPU0:R4(config-bgp-nbr-af)#end
```

To verify the configuration:

On R3

```
R3#Show ip bgp VPNv4 all summary | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.4.4.4         4 65001      12      12        1    0    0 00:08:44        0
```

On R4

```
RP/0/5/CPU0:R4#Sh bgp vpnv4 unicast summary | b Neighbor

Neighbor        Spk    AS MsgRcvd MsgSent   TblVer  InQ OutQ  Up/Down  St/PfxRcd
3.3.3.3           0 65001       5       3        1    0    0 00:01:55          0
```

Task 4

Configure a **Virtual Routing Forwarding (VRF) Instance with a name of CA** (For Customer A), a **route-distinguisher (RD) of "1:10"** and a **route-target (RT) of "1:100"** on R3. On R4, the same route targets will be used for the vrf, but the **RD should be configured to be "1:20"** and the name of the VRF should be configured as **CB**.

On R3

```
R3(config)#ip vrf CA
R3(config-vrf)#RD 1:10
R3(config-vrf)#route-target Both 1:100

R3(config)#int s1/0
R3(config-if)#ip vrf forwarding CA
R3(config-if)#ip address 10.1.13.3 255.255.255.0
```

On R4

```
RP/0/5/CPU0:R4(config)#vrf CB
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config-vrf)#address-family ipv4 uni
RP/0/5/CPU0:R4(config-vrf-af)#import route-target 1:100
RP/0/5/CPU0:R4(config-vrf-af)#export route-target 1:100

RP/0/5/CPU0:R4(config-vrf-af)#int g0/1/0/2
RP/0/5/CPU0:R4(config-if)#no ipv4 address
RP/0/5/CPU0:R4(config-if)#vrf CB
RP/0/5/CPU0:R4(config-if)#ipv4 address 10.1.45.4/24
RP/0/5/CPU0:R4(config-if)#commit
RP/0/5/CPU0:R4(config-if)#end
```

To verify the VRF Configuration:

On R4

```
RP/0/5/CPU0:R4#Sh vrf CB detail

VRF CB; RD not set; VPN ID not set
VRF mode: Regular
Description not set
Interfaces:
  GigabitEthernet0/1/0/2
Address family IPV4 Unicast
  Import VPN route-target communities:
    RT:1:100
  Export VPN route-target communities:
    RT:1:100
  No import route policy
  No export route policy
Address family IPV6 Unicast
  No import VPN route-target communities
  No export VPN route-target communities
  No import route policy
  No export route policy
```

On R3

```
R3#Sh ip vrf detail

VRF CA (VRF Id = 1); default RD 1:10; default VPNID <not set>
  Interfaces:
    Se1/0
VRF Table ID = 1
  Export VPN route-target communities
    RT:1:100
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Import VPN route-target communities
  RT:1:100
No import route-map
No export route-map
VRF label distribution protocol: not configured
VRF label allocation mode: per-prefix
```

Task 5

Configure BGP as the MPLS Routing context between the CE and their respective PE (R1 and R3 as well as R4 and R5). The customer AS of 65015 should be assigned to both customer sites, whereas, the provider should use AS 65001.

Configuring the CE routers:

On R1

```
R1(config)#router bgp 65015
R1(config-router)#No au
R1(config-router)#neigh 10.1.13.3 remot 65001
R1(config-router)#Netw 1.0.0.0
R1(config-router)#Netw 192.168.1.0
```

On R5

```
R5(config)#router bgp 65015
R5(config-router)#No au
R5(config-router)#Neigh 10.1.45.4 remot 65001
R5(config-router)#Netw 5.0.0.0
R5(config-router)#Netw 192.168.5.0
```

Configuring the PE routers to establish a peer session with the CE routers:

Note, no redistribution is necessary for MP-BGP. The peering is identical to that of Normal BGP except it is in the VPN instead. However, normal BGP rules apply within the VPN.

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#rd 1:20
RP/0/5/CPU0:R4(config-bgp-vrf)#address-family ipv4 uni
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config-bgp-vrf-af)#commit
RP/0/5/CPU0:R4(config-bgp-vrf-af)#exit
RP/0/5/CPU0:R4(config-bgp-vrf)#neighbor 10.1.45.5
RP/0/5/CPU0:R4(config-bgp-vrf-nbr)#remote-as 65015
RP/0/5/CPU0:R4(config-bgp-vrf-nbr)#address-family ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-nbr-af)#commit
RP/0/5/CPU0:R4(config-bgp-vrf-nbr-af)#end
```

On R3

```
R3(config)#router bgp 65001
R3(config-router)#No au
R3(config-router)#address-family ipv4 vrf CA
R3(config-router-af)#Neighbor 10.1.13.1 remot 65015
```

You should see the following console messages before moving on:

```
%BGP-5-ADJCHANGE: neighbor 10.1.13.1 vpn vrf CA Up
```

To verify the configuration

On R5

```
R5#Sh ip bgp summ | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.45.4       4 65001      36      36        3    0    0 00:04:48        0
```

On R1

```
R1#Show ip bgp summ | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.13.3       4 65001      51      51        3    0    0 00:03:31        0
```

On R4

```
RP/0/5/CPU0:R4(config)#route-policy ALLOW-ALL
RP/0/5/CPU0:R4(config-rpl)#pass
RP/0/5/CPU0:R4(config-rpl)#end
RP/0/5/CPU0:R4(config)#commit

RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CB
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R4(config-bgp-vrf)#neighbor 10.1.45.5
RP/0/5/CPU0:R4(config-bgp-vrf-nbr)#address-family ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-nbr-af)#route-policy ALLOW-ALL in
RP/0/5/CPU0:R4(config-bgp-vrf-nbr-af)#commit
```

IOS-XR Routers do not receive ALL BGP routes like IOS router do. Since route-maps are not available in IOS-XR, a simple Route-Policy must be created using Route Policy Language (RPL) that permits all routes to PASS.

On R5

```
R5#Sh ip bgp summ | B Neighbor


Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.45.4       4 65001      36      36        3    0    0 00:04:48        0
```

On R1

```
R1#Show ip bgp summ | B Neighbor

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.13.3       4 65001      51      51        3    0    0 00:03:31        0
```

On R3

```
R3#Sh ip bgp vpnv4 vrf CA | B Net

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:10 (default for vrf CA)
*> 1.0.0.0          10.1.13.1                0             0 65015 i
*>i5.0.0.0          4.4.4.4                  0    100      0 65015 i
*> 192.168.1.0      10.1.13.1                0             0 65015 i
*>i192.168.5.0      4.4.4.4                  0    100      0 65015 i
```

On R4

```
RP/0/5/CPU0:R4#Show bgp vrf CB | b Netw

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:20 (default for vrf CB)
*>i1.0.0.0/8        3.3.3.3                  0    100      0 65015 i
*> 5.0.0.0/8        10.1.45.5                0             0 65015 i
*>i192.168.1.0/24   3.3.3.3                  0    100      0 65015 i
*> 192.168.5.0/24   10.1.45.5                0             0 65015 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
R1#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
*> 192.168.1.0      0.0.0.0                  0         32768 i
```

On R5

```
R5#Show ip bgp | B Net

   Network          Next Hop            Metric LocPrf Weight Path
*> 5.0.0.0          0.0.0.0                  0         32768 i
*> 192.168.5.0      0.0.0.0                  0         32768 i
```

Although the routes were verified to have been received via MP-BGP and redistribution is not a problem that can exist here, the routes are not learned by the CE routers. This is a result of a normal BGP process of loop prevention. There are two ways to resolve this matter:

3. **Configuring the CE solution**
4. **Configuring the PE solution**

Configuring the CE routers to resolve the problem:

On R1

```
R1(config)#router bgp 65015
R1(config-router)#neighbor 10.1.13.3 allowas-in
```

This feature can be used as an alternative if the Service Provider controls the CE router/s.
When network 5.0.0.0 /8 is advertised by the remote office (R5) the AS number is prepended (65015), which is used as a loop prevention mechanism in BGP. The prefix is received by the provider and then, advertised to R1, which is in AS 65015 (Same as R5's AS), when R1 receives the route, it sees its own AS number in the AS-Path attribute and it discards the route as a loop avoidance mechanism, the "allowas-in" command allows the local router to receive routes even if the local router sees its own AS in the AS-Path attribute.

To verify the configuration:

On R1

Note R5 ignored its own AS number in the inbound updates

```
R1#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          0.0.0.0                  0         32768 i
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
*> 5.0.0.0            10.1.13.3                              0 65001 65015 i
*> 192.168.1.0        0.0.0.0                  0         32768 i
*> 192.168.5.0        10.1.13.3                              0 65001 65015 i
```

The same feature must be configured on R5:

On R4

```
RP/0/5/CPU0:R4(config)#router bgp 65001
RP/0/5/CPU0:R4(config-bgp)#vrf CB
RP/0/5/CPU0:R4(config-bgp-vrf)#neighbor 10.1.45.5
RP/0/5/CPU0:R4(config-bgp-vrf-nbr)#address-family ipv4 uni
RP/0/5/CPU0:R4(config-bgp-vrf-nbr-af)#as-override
RP/0/5/CPU0:R4(config-bgp-vrf-nbr-af)#commit
```

To verify the configuration:

On R5

```
R5#Show ip bgp | B Network

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.0.0.0          10.1.45.4                              0 65001 65001 i
*> 5.0.0.0          0.0.0.0                  0         32768 i
*> 192.168.1.0      10.1.45.4                              0 65001 65001 i
*> 192.168.5.0      0.0.0.0                  0         32768 i
```

Unlike the previous solution where the AS number was ignored, using this feature, the provider replaces the customer's AS number with its own. The customer MUST allow prepended AS numbers or else they will NOT get the BGP routes. Because from the customers perspective, it looks like the provider prepended its own AS number.

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Advanced MPLS

Implement, Optimize & Troubleshoot

IOS-XR

Traffic-Engineering

www.MicronicsTraining.com

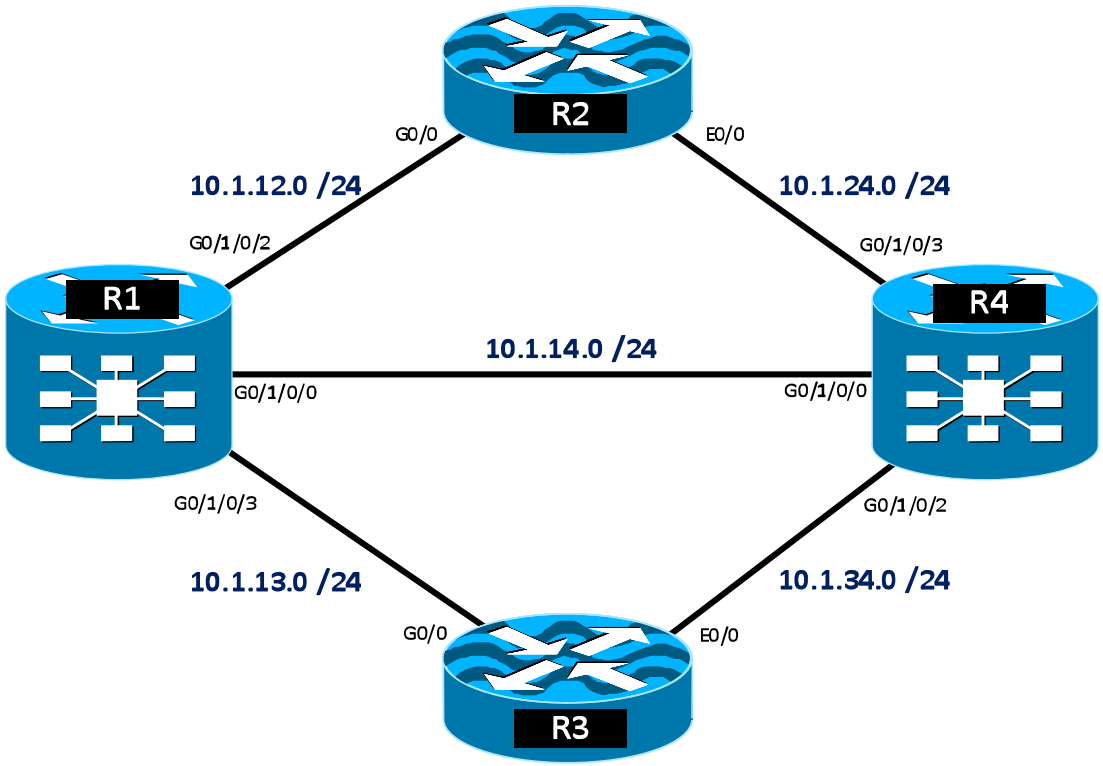
Narbik Kocharians
CCIE #12410
R&S, Security, SP

Paul Negron
CCIE #14856
SP

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Lab 7 – Configuring Static Tunnels



IP Addressing:

Router	Interface	IP address
R1	Lo0 G0/1/0/0 G0/1/0/2 G0/1/0/3	1.1.1.1 /32 10.1.14.1 /24 10.1.12.1 /24 10.1.13.1/ 24
R2	Lo0 G0/0 E0/0	2.2.2.2 /32 10.1.12.2 /24 10.1.24.2 /24
R3	Lo0 G0/0 E0/0	3.3.3.3/ 32 10.1.13.3 /24 10.1.34.3 /24
R4	Lo0 G0/1/0/0	4.4.4.4 /32 10.1.14.4/ 24

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

	G0/1/0/2 G0/1/0/3	10.1.24.4 /24 10.1.34.4 /24

Lab Setup:

XXXXXXXXXXX

Task 1

Configure an OSPF routed core using process "1". Support mpls traffic engineering throughout the core topology.

Configure basic traffic Engineering on ALL Core Devices:

On R1

```
RP/0/5/CPU0:R1(config)#router ospf 1
RP/0/5/CPU0:R1(config-ospf)#area 0
RP/0/5/CPU0:R1(config-ospf-ar)#mpls traffic-eng
RP/0/5/CPU0:R1(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R1(config-ospf-ar-if)#int g0/1/0/2
RP/0/5/CPU0:R1(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R1(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R1(config-ospf-ar-if)#passive
RP/0/5/CPU0:R1(config-ospf-ar-if)#commit
RP/0/5/CPU0:R1(config-ospf-ar-if)#root
```

OSPF needs to recognize and use the OSPF Opaque LSA (type 10) to include the Traffic Engineering paths that will be configured. This command must be configured on ALL core routers.

```
RP/0/5/CPU0:R1(config)#mpls traffic-eng
RP/0/5/CPU0:R1(config-mpls-te)#int g0/1/0/0
RP/0/5/CPU0:R1(config-mpls-te-if)#int g0/1/0/2
RP/0/5/CPU0:R1(config-mpls-te-if)#int g0/1/0/3
RP/0/5/CPU0:R1(config-mpls-te-if)#commit
RP/0/5/CPU0:R1(config-mpls-te-if)#root

RP/0/5/CPU0:R1(config)#rsvp
RP/0/5/CPU0:R1(config-rsvp)#bandwidth rdm max-reservable-bc0 percentage 90
RP/0/5/CPU0:R1(config-rsvp)#int g0/1/0/0
RP/0/5/CPU0:R1(config-rsvp-if)#int g0/1/0/2
RP/0/5/CPU0:R1(config-rsvp-if)#int g0/1/0/3
RP/0/5/CPU0:R1(config-rsvp-if)#commit
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R1(config-rsvp-if)#end
```

ALL interfaces that may carry tunnel traffic at any time must be configured with the "mpls traffic-eng tunnels" command on the interface as well. If this command is not issued at this level, the router will not participate in MPLS Traffic Engineering path calculation (PCALC).

ALL interfaces that may carry tunnel traffic at any time must be configured with the "ip rsvp bandwidth" command on the interface that could be a potential path. The router will not participate in MPLS Traffic Engineering path calculation (PCALC) from the head-end of the tunnel if not configured. The default value, if not specified, will be 75% of the "bandwidth" command configured on the interface. RSVP will use this value to subtract any requested bandwidth by a tunnel.

The three configurations: "mpls traffic-eng tunnels"(global) in IOS only, "mpls traffic-eng tunnels"(interface), and "rsvp"(interface) must all be used on all core interfaces to participate in MPLS Traffic Engineering, even if there will be no tunnels built from the router itself.

On R2

```
R2(config)#mpls traffic-eng tunnels

R2(config)#router ospf 1
R2(config)#network 0.0.0.0 0.0.0.0 area 0
R2(config-router)#mpls traffic-eng area 0
R2(config-router)#mpls traffic-eng router-id loopback 0

R2(config)#interface g0/0
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth

R2(config-if)#interface e0/0
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth
```

On R3

```
R3(config)#mpls traffic-eng tunnels

R3(config)#router ospf 1
R3(config)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#mpls traffic-eng area 0
R3(config-router)#mpls traffic-eng router-id loopback 0

R3(config)# interface g0/0
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-if)#interface e0/0
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 1
RP/0/5/CPU0:R4(config-ospf)#area 0
RP/0/5/CPU0:R4(config-ospf-ar)#mpls traffic-eng
RP/0/5/CPU0:R4(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ospf-ar-if)#int g0/1/0/2
RP/0/5/CPU0:R4(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R4(config-ospf-ar-if)#int loop 0
RP/0/5/CPU0:R4(config-ospf-ar-if)#passive
RP/0/5/CPU0:R4(config-ospf-ar-if)#commit
RP/0/5/CPU0:R4(config-ospf-ar-if)#root

RP/0/5/CPU0:R4(config)#mpls traffic-eng
RP/0/5/CPU0:R4(config-mpls-te)#int g0/1/0/0
RP/0/5/CPU0:R4(config-mpls-te-if)#int g0/1/0/2
RP/0/5/CPU0:R4(config-mpls-te-if)#int g0/1/0/3
RP/0/5/CPU0:R4(config-mpls-te-if)#commit
RP/0/5/CPU0:R4(config-mpls-te-if)#root

RP/0/5/CPU0:R4(config)#rsvp
RP/0/5/CPU0:R4(config-rsvp)#int g0/1/0/0
RP/0/5/CPU0:R4(config-rsvp-if)#int g0/1/0/2
RP/0/5/CPU0:R4(config-rsvp-if)#int g0/1/0/3
RP/0/5/CPU0:R4(config-rsvp-if)#commit
RP/0/5/CPU0:R4(config-rsvp-if)#end
```

Verifying basic traffic Engineering:

On R1

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng link-management igp-neighbors

  Link ID:: GigabitEthernet0/1/0/0
    Neighbor ID: 10.1.14.1 (OSPF 1 area 0, link address: 0.0.0.0)

  Link ID:: GigabitEthernet0/1/0/2
    Neighbor ID: 10.1.12.1 (OSPF 1 area 0, link address: 0.0.0.0)

  Link ID:: GigabitEthernet0/1/0/3
    Neighbor ID: 10.1.13.2 (OSPF 1 area 0, link address: 0.0.0.0)
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

MPLS traffic engineering has been correctly configured under the core interfaces or we would not be able to see neighbors.

```
RP/0/5/CPU0:R1#Sh ospf database | b Type-10

                Type-10 Opaque Link Area Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Opaque
ID
1.0.0.0         1.1.1.1         1730        0x80000023 0x0014f3         0
1.0.0.0         2.2.2.2         1789        0x80000028 0x000eec         0
1.0.0.0         3.3.3.3         1880        0x80000024 0x001adc         0
1.0.0.0         4.4.4.4         1834        0x80000024 0x001ed0         0
1.0.0.1         2.2.2.2         1533        0x80000025 0x008f9e         1
1.0.0.1         3.3.3.3         1880        0x80000026 0x009979         1
1.0.0.2         2.2.2.2         1789        0x80000024 0x006650         2
1.0.0.2         3.3.3.3         1880        0x80000026 0x00aa02         2
1.0.0.9         1.1.1.1         251         0x80000002 0x005121         9
1.0.0.9         4.4.4.4         1834        0x80000026 0x000936         9
1.0.0.11        1.1.1.1         251         0x80000002 0x00d89b        11
1.0.0.11        4.4.4.4         1834        0x80000026 0x000127        11
1.0.0.12        1.1.1.1         251         0x80000002 0x00353a        12
1.0.0.12        4.4.4.4         1834        0x80000026 0x00030f        12
```

The type 10 opaque LSA is used to identify all of the LSA's learned from the MPLS traffic-engineering topology. The highlighted paths all share the same link ID's which represents the tunnel path and all of the LSA's learned from that path.

```
RP/0/5/CPU0:R1#Sh rsvp int

*: RDM: Default I/F B/W % : 90% [cfgd] (max resv/bc0), 0% [default]
(bc1)

Interface   MaxBW (bps)  MaxFlow (bps) Allocated (bps)      MaxSub (bps)
----------- ------------ ------------- -------------------- ------------
-
Gi0/1/0/0             0             0             0 (  0%)            0
Gi0/1/0/2             0             0             0 (  0%)            0
Gi0/1/0/3             0             0             0 (  0%)            0
```

RSVP has been configured under each of the core interfaces as well. The default-reserved bandwidth is 75% of the physical bandwidth. Bandwidth is a constraint that is not a driving factor in MPLS. The label creation mechanism is what is needed for a successful deployment.

Configure R1 g0/1/0/3 and g0/1/0/2 with "rsvp bandwidth 60" on R1 and R4 before proceeding!! An example is given below:

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
RP/0/5/CPU0:R1(config)#ip rsvp
RP/0/5/CPU0:R1(config-rsvp)#int g0/1/0/3
RP/0/5/CPU0:R1(config-rsvp-if)#bandwidth 60 kbps

The bandwidth can be specified in kbps, Mbps and Gbps.

RP/0/5/CPU0:R1#Sh rsvp int summary

Interface          MaxBW (bps) Allocated (bps) Path In Path Out Resv In Resv Out
------------------ ----------- --------------- ------- -------- ------- --------
Gi0/1/0/0                    0        0 (  0%)       0        0       0        0
Gi0/1/0/2                60000        0 (  0%)       0        1       1        0
Gi0/1/0/3                60000        0 (  0%)       0        0       0        0
```

Task 2

Configure MPLS-TE to establish basic connectivity from R1 to R4. R1 should not follow the default routing path. The path through R2 should be favored as a static path with the following attributes. Priority = 4, all traffic uses the tunnel to 4.4.4.4 and bandwidth = 30K.

Verify Current Routing Path:

On R1

```
RP/0/5/CPU0:R1#Trace 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4

 1  10.1.14.4 112 msec  *  5 msec
 2

R1#Sh ip cef exact-route 1.1.1.1 4.4.4.4

4.4.4.4/32, version 139, internal 0x4000001 (ptr 0xae187010) [1], 0x0
(0xae155298), 0x0 (0x0)
 Updated May 10 04:55:07.557
 local adjacency 10.1.14.4
 Prefix Len 32, traffic index 0, precedence routine (0), priority 1
   via GigabitEthernet0/1/0/0
   via 10.1.14.4, GigabitEthernet0/1/0/0, 5 dependencies, weight 0, class
0 [flags 0x0]
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    path-idx 0 [0xaed35174 0x0]
    next hop 10.1.14.4
    local adjacency
```

The current path is currently R1-R4 through the G0/1/0/0 interface.

Change the Current Routing Path:

```
RP/0/5/CPU0:R1(config)#int tunnel-te 0
RP/0/5/CPU0:R1(config-if)#ipv4 unnumbered lo0
```

In IOS-XR there are separate interface tupes for tunnels instead of tunnel modes. A good practice to not waste an ip address on the tunnel interface is to configure them as unnumbered.

```
RP/0/5/CPU0:R1(config-if)#tunnel destination 4.4.4.4
```

The destination of the tunnel = the traffic-eng router-id of the last hop router

```
RP/0/5/CPU0:R1(config-if)#autoroute announce
```

This command ensures that all traffic is routed into the tunnel and announced into ospf updates. The "tunnel mpls" keyword is not required in IOS-XR.

```
RP/0/5/CPU0:R1(config-if)#priority 4 4
```

The tunnel priority will use the SETUP priority(the first number) and compare with another tunnel's HOLD priority(the second number in the command). Both numbers are usually set to be the same to avoid confusion but may be different to establish true priority.

```
RP/0/5/CPU0:R1(config-if)#signalled-bandwidth 30
RP/0/5/CPU0:R1(config-if)#path-option 1 explicit name SLOW_LINK
RP/0/5/CPU0:R1(config-if)#exit
```

The tunnel signaled-bandwidth of "30" is in kbps and is what the tunnel requires to be setup.
The tunnel should still remain down until the explicit path is created.

```
RP/0/5/CPU0:R1(config)#explicit-path name SLOW_LINK
RP/0/5/CPU0:R1(config-expl-path)#index 1 next-address 10.1.12.2
RP/0/5/CPU0:R1(config-expl-path)#index 2 next-address 10.1.24.4
RP/0/5/CPU0:R1(config-expl-path)#end
```

The path should come up after the last hop is configured

Verify the Traffic Engineered path Information

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels brie

                      TUNNEL NAME          DESTINATION      STATUS  STATE
                       tunnel-te0              4.4.4.4          up  up
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads


RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels

Name: tunnel-te0  Destination: 4.4.4.4
  Status:
    Admin:    up Oper:   up   Path:  valid   Signalling: connected

    path option 1,  type explicit SLOW_LINK (Basis for Setup, path weight
11)
    G-PID: 0x0800 (derived from egress interface properties)
    Bandwidth Requested: 30 kbps  CT0
    Creation Time: Fri May 11 00:43:15 2012 (00:30:46 ago)
  Config Parameters:
    Bandwidth:       30 kbps (CT0) Priority:  4  4 Affinity: 0x0/0xffff
    Metric Type: TE (default)
    Hop-limit: disabled
    AutoRoute:  enabled  LockDown: disabled   Policy class: not set
    Forwarding-Adjacency: disabled
    Loadshare:          0 equal loadshares
    Auto-bw: disabled
    Fast Reroute: Disabled, Protection Desired: None
    Path Protection: Not Enabled
    Soft Preemption: Disabled
  History:
    Tunnel has been up for: 00:05:50 (since Fri May 11 01:08:11 UTC 2012)
    Current LSP:
      Uptime: 00:05:50 (since Fri May 11 01:08:11 UTC 2012)

  Path info (OSPF 1 area 0):
  Hop0: 10.1.12.2
  Hop1: 10.1.24.2
  Hop2: 10.1.24.4
  Hop3: 4.4.4.4
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R1#Sh rsvp int

*: RDM: Default I/F B/W % : 90% [cfgd] (max resv/bc0), 0% [default] (bc1)

Interface   MaxBW (bps)  MaxFlow (bps) Allocated (bps)       MaxSub (bps)
----------- ------------ ------------- -------------------- -------------
Gi0/1/0/0              0             0              0 (  0%)             0
Gi0/1/0/2            60K           60K            30K ( 50%)             0*
Gi0/1/0/3            60K           60K              0 (  0%)             0*
```

The interface through R2 is reserving the bandwidth as 30K.

Control and Data Plane:

On R4

```
RP/0/5/CPU0:R4#Sh mpls traffic-eng  tunnels int inout g0/1/0/2 | i InLabel

  InLabel: GigabitEthernet0/1/0/2, implicit-null
```

R1 sends an RSVP "PATH message" that request for a label concerning the tunnel.
The request passes through R2 and eventually reaches R4. Once the request is received for the tunnel, it is assigned a label of "implicit null" and will send this label to R2. The "OutLabel" represents the action that will be taken in the event it receives the "InLabel" context. In this case, this is the termination point so there should be no Outlabel Context.

On R2

```
R2#Sh mpls traffic-eng tunnels int  | b InLabel |Out

R2#Sh mpls traffic-eng tunnels int  | b InLabel
  InLabel  : GigabitEthernet0/0, 16
  Prev Hop : 10.1.12.1
  OutLabel : Ethernet0/0, implicit-null
  Next Hop : 10.1.24.4
  RSVP Signalling Info:
       Src 1.1.1.1, Dst 4.4.4.4, Tun_Id 0, Tun_Instance 2
    RSVP Path Info:
      My Address: 10.1.24.2
      Explicit Route: 10.1.24.4 4.4.4.4
      Record   Route:   NONE
      Tspec: ave rate=30 kbits, burst=1000 bytes, peak rate=30 kbits
    RSVP Resv Info:
      Record   Route:   NONE
      Fspec: ave rate=30 kbits, burst=1000 bytes, peak rate=30 kbits
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    [LSP Tunnel 1.1.1.1_2->4.4.4.4_0 {7} described, connected]

R2#Sh mpls traffic-eng tun brie | b P2P

P2P TUNNELS/LSPs:
TUNNEL NAME                      DESTINATION      UP IF      DOWN IF
STATE/PROT
R1_t0                            4.4.4.4          Gi0/0      Et0/0
up/up
Displayed 0 (of 0) heads, 1 (of 1) midpoints, 0 (of 0) tails

P2MP TUNNELS:
Displayed 0 (of 0) P2MP heads

P2MP SUB-LSPS:
Displayed 0 P2MP sub-LSPs:
          0 (of 0) heads, 0 (of 0) midpoints, 0 (of 0) tails

R2#Sh mpls forwarding-table | b Local

Local      Outgoing   Prefix           Bytes Label   Outgoing    Next Hop
Label      Label      or Tunnel Id     Switched      interface
16         Pop Label  1.1.1.1 0 [2]    0             Et0/0       10.1.24.4
```

Control Plane: R2 assignes a local tag of "16" for the tunnel when it receives the implicit null tag from R4. The tag could be different in your output.

Data Plane: R2 is waiting for label "16" on the InLabel interface s0/0.21 and forwarding the route with the implicit null tag (label "3") out of "OutLabel" int s0/0.24.

On R1

```
RP/0/5/CPU0:R1#Sh mpls forwarding tunnels

Tunnel    Outgoing    Outgoing     Next Hop         Bytes
Name      Label       Interface                     Switched
--------- ----------- ------------ --------------- ------------
tt0       16          Gi0/1/0/2    10.1.12.2        0
```

Control Plane: R1 learns the "16 label from R2.
Data Plane: Label "16" will be imposed if R1 receives a packet destined to use the tunnel.

```
RP/0/5/CPU0:R1#Trace 4.4.4.4

Type escape sequence to abort.
Tracing the route to 4.4.4.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
 1  10.1.12.2 [MPLS: Label 16 Exp 0] 7 msec  4 msec  4 msec
 2  10.1.24.4 6 msec  *  5 msec

RP/0/5/CPU0:R1#Sh cef exact-route 1.1.1.1 4.4.4.4

4.4.4.4/32, version 145, internal 0x4000001 (ptr 0xae187010) [1], 0x0
(0xae155298), 0x0 (0x0)
 Updated May 11 01:08:11.843
 Prefix Len 32, traffic index 0, precedence routine (0), priority 1
   via tunnel-te0
   via 4.4.4.4, tunnel-te0, 3 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 [0xaed031c8 0x0]
    next hop 4.4.4.4
    local adjacency
```

R1 uses the slower path due to the TE configuration on the R1 tunnel head.

Task 3

Configure a second tunnel (tunnel ”1”) on R1 that uses the same path option configured in Task 2. Provide an additional path option 2 MEDIUM_LINK on both tunnels. This path option should use the path through R3.

On R1

```
RP/0/5/CPU0:R1(config)#int tunnel-te0
RP/0/5/CPU0:R1(config-if)#path-option 2 explicit name MEDIUM_LINK
RP/0/5/CPU0:R1(config-if)#root
```

We add the second path option on tunnel 0

```
RP/0/5/CPU0:R1(config)#int tunnel-te1
RP/0/5/CPU0:R1(config-if)#ipv4 unnumbered lo0
RP/0/5/CPU0:R1(config-if)#destination 4.4.4.4
RP/0/5/CPU0:R1(config-if)#autoroute announce
RP/0/5/CPU0:R1(config-if)#priority 4 4
RP/0/5/CPU0:R1(config-if)#bandwidth 30
RP/0/5/CPU0:R1(config-if)#path-option 1 explicit name SLOW_LINK
RP/0/5/CPU0:R1(config-if)#path-option 2 explicit name MEDIUM_LINK
RP/0/5/CPU0:R1(config-if)#root
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The second tunnel is created with similar options for simplicity. The path will not come up until at least one of the explicit paths is configured.

```
RP/0/5/CPU0:R1(config)#explicit-path name MEDIUM_LINK
RP/0/5/CPU0:R1(config-expl-path)#index 1 next-address 10.1.13.3
RP/0/5/CPU0:R1(config-expl-path)#index 1 next-address 10.1.34.4
RP/0/5/CPU0:R1(config-expl-path)#commit
RP/0/5/CPU0:R1(config-expl-path)#end
```

The tunnel comes up immediately. 2 paths have been created. We need to see what tunnel is being used and analyze the constraints.

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels name tunnel-te0

Name: tunnel-te0  Destination: 4.4.4.4
  Status:
    Admin:    up Oper:   up   Path:  valid   Signalling: connected

    path option 1,  type explicit SLOW_LINK (Basis for Setup, path weight
11)
    path option 2,  type explicit MEDIUM_LINK
    G-PID: 0x0800 (derived from egress interface properties)
    Bandwidth Requested: 30 kbps  CT0
    Creation Time: Fri May 11 00:43:15 2012 (1d02h ago)
  Config Parameters:
    Bandwidth:       30 kbps (CT0) Priority:  4  4 Affinity: 0x0/0xffff
    Metric Type: TE (default)
    Hop-limit: disabled
    AutoRoute:  enabled  LockDown: disabled   Policy class: not set
    Forwarding-Adjacency: disabled
    Loadshare:          0 equal loadshares
    Auto-bw: disabled
    Fast Reroute: Disabled, Protection Desired: None
    Path Protection: Not Enabled
    Soft Preemption: Disabled
  History:
    Tunnel has been up for: 1d02h (since Fri May 11 01:08:11 UTC 2012)
    Current LSP:
      Uptime: 1d02h (since Fri May 11 01:08:11 UTC 2012)
    Reopt. LSP:
      Last Failure:
        LSP not signalled, identical to the [CURRENT] LSP
        Date/Time: Fri May 11 01:23:53 UTC 2012 [1d01h ago]

  Path info (OSPF 1 area 0):
  Hop0: 10.1.12.2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
  Hop1: 10.1.24.2
  Hop2: 10.1.24.4
  Hop3: 4.4.4.4
Displayed 1 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads---text omitted---
```

The "Tunnel 0" interface uses the 1st path option since the tunnel priority does not come into play in this scenario and the bandwidth is available.

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels name tunnel-te1

Name: tunnel-te1  Destination: 4.4.4.4
  Status:
    Admin:    up Oper:   up   Path:  valid   Signalling: connected

    path option 1,  type explicit SLOW_LINK (Basis for Setup, path weight
11)
    path option 2,  type explicit MEDIUM_LINK
    G-PID: 0x0800 (derived from egress interface properties)
    Bandwidth Requested: 0 kbps  CT0
    Creation Time: Sat May 12 03:07:18 2012 (00:03:45 ago)
  Config Parameters:
    Bandwidth:        0 kbps (CT0) Priority:  4  4 Affinity: 0x0/0xffff
    Metric Type: TE (default)
    Hop-limit: disabled
    AutoRoute:  enabled  LockDown: disabled   Policy class: not set
    Forwarding-Adjacency: disabled
    Loadshare:          0 equal loadshares
    Auto-bw: disabled
    Fast Reroute: Disabled, Protection Desired: None
    Path Protection: Not Enabled
    Soft Preemption: Disabled
  History:
    Tunnel has been up for: 00:03:44 (since Sat May 12 03:07:19 UTC 2012)
    Current LSP:
      Uptime: 00:03:44 (since Sat May 12 03:07:19 UTC 2012)

  Path info (OSPF 1 area 0):
  Hop0: 10.1.12.2
  Hop1: 10.1.24.2
  Hop2: 10.1.24.4
  Hop3: 4.4.4.4
Displayed 1 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads
```

"Tunnel 1" also uses the SLOW_LINK. There is enough bandwidth available and it is also the first

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

option on this tunnel as well.

Task 4

Configure Tunnel "1 "to use 40 K of bandwidth and reoptomize the tunnels. Since RSVP bandwidth has been limited to 60K on every interface, view the changes that result in reoptimization of the tunnels.

On R1

```
RP/0/5/CPU0:R1(config)#int tunnel-te1
RP/0/5/CPU0:R1(config-if)#bandwidth 40
RP/0/5/CPU0:R1(config-if)#commit
RP/0/5/CPU0:R1(config-if)#end
```

We will view the effects of changing the bandwidth requirement of "Tunnel 1"

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels brie

                      TUNNEL NAME          DESTINATION      STATUS  STATE
                       tunnel-te0              4.4.4.4          up   up
                       tunnel-te1              4.4.4.4          up   up
Displayed 2 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 2 up, 0 down, 0 recovering, 0 recovered heads
```

The traffic that is currently using the SLOW_LINK may not change until reoptimization takes place (3600 seconds). In ths case, a changing of the resources has taken place, which triggers TE path change due to the constraint. Disabling timer frequency or events will have no effect.

```
RP/0/5/CPU0:R1#Show mpls traffic-eng tunnel name tunnel-te1

Name: tunnel-te1  Destination: 4.4.4.4
  Status:
    Admin:    up Oper:   up   Path:  valid   Signalling: connected

    path option 2,  type explicit MEDIUM_LINK (Basis for Setup, path
weight 11)
    path option 1,  type explicit SLOW_LINK
    Last PCALC Error [Reopt]: Sat May 12 03:21:47 2012
      Info: Can't reach 10.1.12.2 on 2.2.2.2, from node 1.1.1.1 (bw)
    G-PID: 0x0800 (derived from egress interface properties)
    Bandwidth Requested: 40 kbps  CT0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    Creation Time: Sat May 12 03:07:18 2012 (00:15:08 ago)
  Config Parameters:
    Bandwidth:        40 kbps (CT0) Priority:  4  4 Affinity: 0x0/0xffff
    Metric Type: TE (default)
    Hop-limit: disabled
    AutoRoute:  enabled  LockDown: disabled   Policy class: not set
    Forwarding-Adjacency: disabled
    Loadshare:           0 equal loadshares
    Auto-bw: disabled
    Fast Reroute: Disabled, Protection Desired: None
    Path Protection: Not Enabled
    Soft Preemption: Disabled
  History:
    Tunnel has been up for: 00:15:07 (since Sat May 12 03:07:19 UTC 2012)
    Current LSP:
      Uptime: 00:00:39 (since Sat May 12 03:21:47 UTC 2012)
    Reopt. LSP:
      Last Failure:
        LSP not signalled, identical to the [CURRENT] LSP
        Date/Time: Sat May 12 03:21:02 UTC 2012 [00:01:24 ago]
    Prior LSP:
      ID: path option 1 [2]
      Removal Trigger: reoptimization completed

  Path info (OSPF 1 area 0):
  Hop0: 10.1.13.3
  Hop1: 10.1.34.3
  Hop2: 10.1.34.4
  Hop3: 4.4.4.4
Displayed 1 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads
```

Notice how although the SLOW_LINK is the 1st option, the MEDIUM_PATH is chosen because the primary link only has a bandwidth of 60 kbps configured on the interface. The "Tunnel 0" interface has a requirement of 30 kbps. The Tunnel 1" interface has a requirement of 40 kbps and RSVP will reject the reservation on the primary path. The second path is unused and is chosen as the best for this tunnel.
Also, the Current LSP information displays that the new LSP was due to reoptimization. The reoptimization is now automatic when a change is performed on the tunnel and does not necessarily need to be manually triggered. In Cisco IOS-XR, The lowered number tunnel will win if ALL constraints are tied except for bandwidth.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 5

Configure the tunnels such that when shutdown manually, Tunnel"1" will use the primary path no matter which order the tunnels come up. Even during reoptimization.

On R1

```
RP/0/5/CPU0:R1(config)#int tunnel-te0
RP/0/5/CPU0:R1(config-if)#shut
RP/0/5/CPU0:R1(config-if)#int tunnel-te1
RP/0/5/CPU0:R1(config-if)#shut
RP/0/5/CPU0:R1(config-if)#commit
RP/0/5/CPU0:R1(config-if)#root

RP/0/5/CPU0:R1(config-if)#priority 3 4
RP/0/5/CPU0:R1(config-if)#commit

% Failed to commit one or more configuration items during a pseudo-atomic
operation. All changes made have been reverted. Please issue 'show
configuration failed' from this session to view the errors

RP/0/5/CPU0:R1(config-if)#root

RP/0/5/CPU0:R1(config)#Show config failed

!! SEMANTIC ERRORS: This configuration was rejected by
!! the system due to semantic errors. The individual
!! errors with each failed configuration command can be
!! found below.

interface tunnel-te1
 priority 3 4
!!% Invalid or corrupt data was supplied to TE: Setup priority may not be
higher than hold priority
!
end

RP/0/5/CPU0:R1(config)#int tunnel-te1
RP/0/5/CPU0:R1(config-if)#priority 3 3
RP/0/5/CPU0:R1(config-if)#commit
RP/0/5/CPU0:R1(config-if)#root
```

The priority is used to accomplish the task correctly. The hold priority cannot be a higher number (inferior) than the setup priority. They are normally set to be identical but the hold priority can be set lower (better) than the setup priority. In some cases the traffic using the tunnel may not need a better

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

priority to be established but once it is established it should not be taken down.

```
RP/0/5/CPU0:R1(config)#int tunnel-te0
RP/0/5/CPU0:R1(config-if)#NO shut
RP/0/5/CPU0:R1(config-if)#commit
RP/0/5/CPU0:R1(config-if)#end

RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnel brie | b TUNNEL

                      TUNNEL NAME           DESTINATION       STATUS  STATE
                       tunnel-te0             4.4.4.4           up    up
                       tunnel-te1             4.4.4.4         down    down
Displayed 2 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 1 down, 0 recovering, 0 recovered head
```

We bring up tunnel "0" first and make sure it is established. It uses the primary path through R2.

```
RP/0/5/CPU0:R1(config)#int tunnel-te1
RP/0/5/CPU0:R1(config-if)#NO shut
RP/0/5/CPU0:R1(config-if)#commit
RP/0/5/CPU0:R1(config-if)#end
```

Verifying the Priority constraint:

On R1

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels name tunnel-te1 | i path option

    path option 1,  type explicit SLOW_LINK (Basis for Setup, path weight
11)
    path option 2,  type explicit MEDIUM_LINK
```

As we can see. Tunnel "1" gets established along the primary path. It bumps off Tunnel 0" in due to the better setup priority.

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels name tunnel-te0 | i path option

    path option 2,  type explicit MEDIUM_LINK (Basis for Setup, path
weight 11)
    path option 1,  type explicit SLOW_LINK
      ID: path option 1 [8]
```

Tunnel "0" is setup across the MEDIUM_LINK path instead.

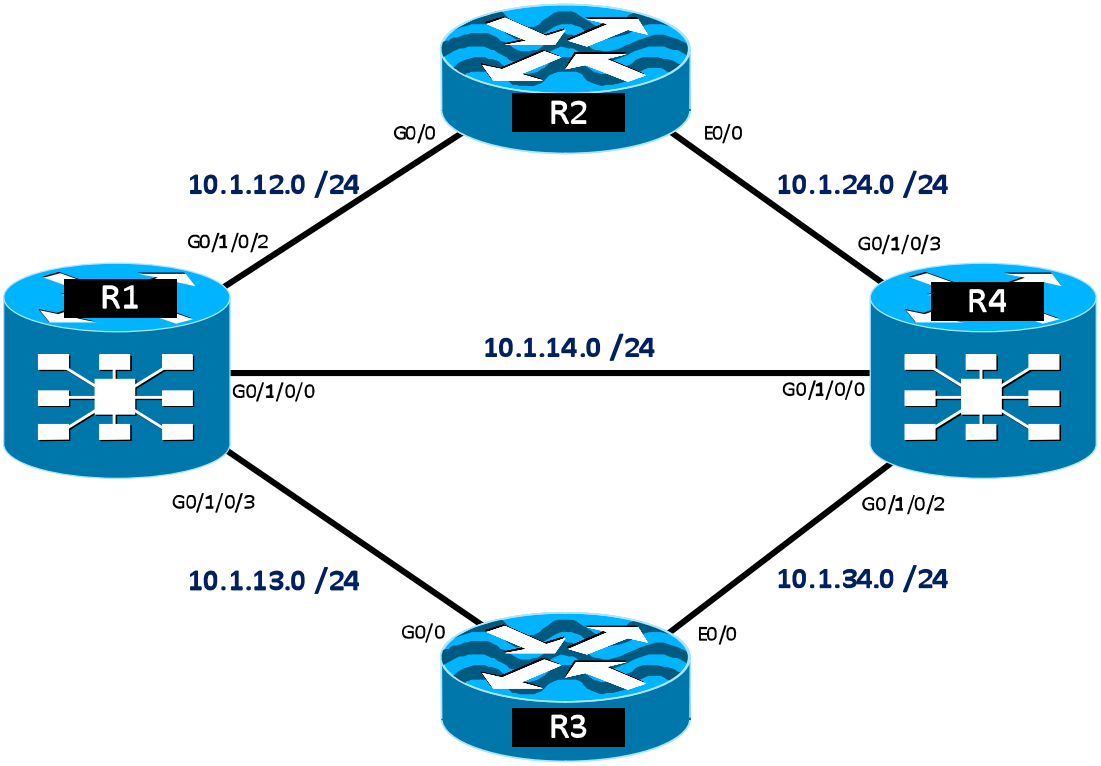
©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 6

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Lab 8 – Dynamic Tunnels



IP Addressing:

Router	Interface	IP address
R1	Lo0 G0/1/0/0 G0/1/0/2 G0/1/0/3	1.1.1.1 /32 10.1.14.1 /24 10.1.12.1 /24 10.1.13.1/ 24
R2	Lo0	2.2.2.2 /32

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

	G0/0 E0/0	10.1.12.2 /24 10.1.24.2 /24
R3	Lo0 G0/0 E0/0	3.3.3.3/ 32 10.1.13.3 /24 10.1.34.3 /24
R4	Lo0 G0/1/0/0 G0/1/0/2 G0/1/0/3	4.4.4.4 /32 44.4.4.4/32 10.1.14.4/ 24 10.1.24.4 /24 10.1.34.4 /24

Lab Setup:

The lab must have basic IGP and MPLS traffic Engineering configured to support tunnels. RSVP should use a bandwidth of 80 Kbps on all paths except for R1>R3>R4.

Task 1

Configure an OSPF with a process of "1" on all routers. The routers must be prepared to support MPLS-TE. Complete IGP connectivity between the devices is required. AN OSPF cost of "130" shoud be used on path R1>R3>R4. All other paths should have an ospf cost of "80".

On R1

```
RP/0/5/CPU0:R1(config)#router ospf 1
RP/0/5/CPU0:R1(config-ospf)#mpls traffic-eng router-id lo0
RP/0/5/CPU0:R1(config-ospf)#area 0
RP/0/5/CPU0:R1(config-ospf)#cost 64
RP/0/5/CPU0:R1(config-ospf-ar)#mpls traffic-eng
RP/0/5/CPU0:R1(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R1(config-ospf-ar-if)#int g0/1/0/2
RP/0/5/CPU0:R1(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R1(config-ospf-ar-if)#cost 130
RP/0/5/CPU0:R1(config-ospf-ar-if)#int lo0
RP/0/5/CPU0:R1(config-ospf-ar-if)#passive
RP/0/5/CPU0:R1(config-ospf-ar-if)#commit
RP/0/5/CPU0:R1(config-ospf-ar-if)#root

RP/0/5/CPU0:R1(config)#mpls traffic-eng
RP/0/5/CPU0:R1(config-mpls-te)#int g0/1/0/0
RP/0/5/CPU0:R1(config-mpls-te-if)#int g0/1/0/2
RP/0/5/CPU0:R1(config-mpls-te-if)#int g0/1/0/3
RP/0/5/CPU0:R1(config-mpls-te-if)#commit
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R1(config-mpls-te-if)#root

RP/0/5/CPU0:R1(config)#rsvp
RP/0/5/CPU0:R1(config-rsvp)#int g0/1/0/2
RP/0/5/CPU0:R1(config-rsvp-if)#bandwidth 80
RP/0/5/CPU0:R1(config-rsvp-if)#int g0/1/0/0
RP/0/5/CPU0:R1(config-rsvp-if)#bandwidth 80
RP/0/5/CPU0:R1(config-rsvp-if)#int g0/1/0/3
RP/0/5/CPU0:R1(config-rsvp-if)#bandwidth 768
RP/0/5/CPU0:R1(config-rsvp-if)#commit
RP/0/5/CPU0:R1(config-mpls-te-if)#end
```

On R2

```
R2(config)#mpls traffic-eng tunnels

R2(config)#router ospf 1
R2(config)#network 0.0.0.0 0.0.0.0 area 0
R2(config-router)#mpls traffic-eng area 0
R2(config-router)#mpls traffic-eng router-id loopback 0

R2(config)# interface g0/0
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth 80
R2(config-if)#ip ospf cost 80

R2(config-if)#interface e0/0
R2(config-if)#mpls traffic-eng tunnels
R2(config-if)#ip rsvp bandwidth 80
R2(config-if)#ip ospf cost 80
```

On R3

```
R3(config)#mpls traffic-eng tunnels

R3(config)#router ospf 1
R3(config)#network 0.0.0.0 0.0.0.0 area 0
R3(config-router)#mpls traffic-eng area 0
R3(config-router)#mpls traffic-eng router-id loopback 0

R3(config)# interface g0/0
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth 768
R3(config-if)#ip ospf cost 768

R3(config-if)#interface e0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3(config-if)#mpls traffic-eng tunnels
R3(config-if)#ip rsvp bandwidth 768
R2(config-if)#ip ospf cost 768
```

On R4

```
RP/0/5/CPU0:R4(config)#router ospf 1
RP/0/5/CPU0:R4(config-ospf)#mpls traffic-eng router-id lo0
RP/0/5/CPU0:R4(config-ospf)#area 0
RP/0/5/CPU0:R4(config-ospf-ar)#cost 64
RP/0/5/CPU0:R4(config-ospf-ar)#mpls traffic-eng
RP/0/5/CPU0:R4(config-ospf-ar)#int g0/1/0/0
RP/0/5/CPU0:R4(config-ospf-ar-if)#int g0/1/0/2
RP/0/5/CPU0:R4(config-ospf-ar-if)#int g0/1/0/3
RP/0/5/CPU0:R4(config-ospf-ar-if)#cost 130
RP/0/5/CPU0:R4(config-ospf-ar-if)#lo0
RP/0/5/CPU0:R4(config-ospf-ar-if)#passive
RP/0/5/CPU0:R4(config-ospf-ar-if)#lo1
RP/0/5/CPU0:R4(config-ospf-ar-if)#passive
RP/0/5/CPU0:R4(config-ospf-ar-if)#commit
RP/0/5/CPU0:R4(config-ospf-ar-if)#root

RP/0/5/CPU0:R4(config)#mpls traffic-eng
RP/0/5/CPU0:R4(config-mpls-te)#int g0/1/0/0
RP/0/5/CPU0:R4(config-mpls-te-if)#int g0/1/0/2
RP/0/5/CPU0:R4(config-mpls-te-if)#int g0/1/0/3
RP/0/5/CPU0:R4(config-mpls-te-if)#commit
RP/0/5/CPU0:R4(config-mpls-te-if)#root

RP/0/5/CPU0:R4(config)#rsvp
RP/0/5/CPU0:R4(config-rsvp)#int g0/1/0/2
RP/0/5/CPU0:R4(config-rsvp-if)#bandwidth 80
RP/0/5/CPU0:R4(config-rsvp-if)#int g0/1/0/0
RP/0/5/CPU0:R4(config-rsvp-if)#bandwidth 80
RP/0/5/CPU0:R4(config-rsvp-if)#int g0/1/0/3
RP/0/5/CPU0:R4(config-rsvp-if)#bandwidth 768
RP/0/5/CPU0:R4(config-rsvp-if)#commit
RP/0/5/CPU0:R4(config-rsvp-if)#end
```

Verify the Configuration

On R1

```
RP/0/5/CPU0:R1#Sh route ospf | i O
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
O    2.2.2.2/32 [110/65] via 10.1.12.2, 00:15:58, GigabitEthernet0/1/0/2
O    3.3.3.3/32 [110/131] via 10.1.13.3, 00:13:26, GigabitEthernet0/1/0/3
O    4.4.4.4/32 [110/128] via 10.1.14.4, 00:08:08, GigabitEthernet0/1/0/0
O    10.1.24.0/24 [110/128] via 10.1.14.4, 00:07:56, GigabitEthernet0/1/0/0
O    10.1.34.0/24 [110/194] via 10.1.14.4, 00:06:10, GigabitEthernet0/1/0/0
O    44.4.4.4/32 [110/128] via 10.1.14.4, 00:08:08, GigabitEthernet0/1/0/0

RP/0/5/CPU0:R1#Show mpls int

Interface                  LDP      Tunnel   Enabled
-------------------------- -------- -------- --------
GigabitEthernet0/1/0/0     No       Yes      Yes
GigabitEthernet0/1/0/2     No       Yes      Yes
GigabitEthernet0/1/0/3     No       Yes      Yes
```

On R2

```
R2#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
O        1.1.1.1 [110/144] via 10.1.12.1, 00:08:59, GigabitEthernet0/0
O        3.3.3.3 [110/211] via 10.1.24.4, 00:08:59, Ethernet0/0
O        4.4.4.4 [110/144] via 10.1.24.4, 00:08:59, Ethernet0/0
O        10.1.13.0/24 [110/210] via 10.1.12.1, 00:08:59, GigabitEthernet0/0
O        10.1.14.0/24 [110/144] via 10.1.24.4, 00:08:59, Ethernet0/0
O        10.1.34.0/24 [110/210] via 10.1.24.4, 00:08:59, Ethernet0/0
O        44.4.4.4 [110/144] via 10.1.24.4, 00:08:59, Ethernet0/0

R2#Show mpls int

Interface              IP             Tunnel   BGP Static Operational
Ethernet0/0            No             Yes      No  No     Yes
GigabitEthernet0/0     No             Yes      No  No     Yes
```

The Tunnel Section is specifically for the TE configuration.

On R3

```
R3#Sh ip route ospf | i O

       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
O        1.1.1.1 [110/832] via 10.1.13.2, 00:08:09, GigabitEthernet0/0
O        2.2.2.2 [110/833] via 10.1.34.4, 00:08:09, Ethernet0/0
O        4.4.4.4 [110/832] via 10.1.34.4, 00:08:09, Ethernet0/0
O        10.1.12.0/24 [110/832] via 10.1.13.2, 00:08:09, GigabitEthernet0/0
O        10.1.14.0/24 [110/832] via 10.1.34.4, 00:08:09, Ethernet0/0
O        10.1.24.0/24 [110/832] via 10.1.34.4, 00:08:09, Ethernet0/0
O        44.4.4.4 [110/832] via 10.1.34.4, 00:08:09, Ethernet0/0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
R3#Sh mpls int

Interface              IP             Tunnel   BGP Static Operational
Ethernet0/0            No             Yes      No  No     Yes
GigabitEthernet0/0     No             Yes      No  No     Yes
```

On R4

```
RP/0/5/CPU0:R4#Sh route ospf | i O

O    1.1.1.1/32 [110/128] via 10.1.14.1, 00:10:58, GigabitEthernet0/1/0/0
O    2.2.2.2/32 [110/65] via 10.1.24.2, 00:10:53, GigabitEthernet0/1/0/2
O    3.3.3.3/32 [110/131] via 10.1.34.3, 00:11:02, GigabitEthernet0/1/0/3
O    10.1.12.0/24 [110/128] via 10.1.14.1, 00:10:58, GigabitEthernet0/1/0/0
O    10.1.13.0/24 [110/194] via 10.1.14.1, 00:09:08, GigabitEthernet0/1/0/0

RP/0/5/CPU0:R4#Sh mpls int

Interface                  LDP      Tunnel   Enabled
-------------------------- -------- -------- --------
GigabitEthernet0/1/0/0     No       Yes      Yes
GigabitEthernet0/1/0/2     No       Yes      Yes
GigabitEthernet0/1/0/3     No       Yes      Yes
```

Task 2

Create a tunnel from R1 to R4 (Tunnel0) that automatically builds according to the current cost of the links. The signalled bandwidth should be 50K. Tunnel Priority should be setup of 1 and hold of 1. All traffic from any source that terminates to 4.4.4.4 should use the tunnel.

On R1

```
RP/0/5/CPU0:R1(config)#int tunnel-te0
RP/0/5/CPU0:R1(config-if)#ipv4 unnumbered lo0
RP/0/5/CPU0:R1(config-if)#destination 4.4.4.4
RP/0/5/CPU0:R1(config-if)#autoroute announce
RP/0/5/CPU0:R1(config-if)#priority 1 1
RP/0/5/CPU0:R1(config-if)#signalled-bandwidth 50
RP/0/5/CPU0:R1(config-if)#path-option 1 dynamic
```

The path option uses the "dynamic" keyword instead of static to accomplish the task.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Verify the dynamic tunnel:

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels up | b Path info

  Path info (OSPF 1 area 0):
  Node hop count: 1
  Hop0: 10.1.14.4
  Hop1: 4.4.4.4
Displayed 1 (of 1) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads
```

The path direct to R4 is chosen as the best path with the necessary bandwidth available.

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng link-management admission-control int g0/1/0/0

  System Information::
      Tunnels Count     : 1
      Tunnels Selected  : 1

  Bandwidth descriptor legend:
    B0 = bw from pool 0, B1 = bw from pool 1, R = bw locked, H = bw held

TUNNEL ID                UP IF      DOWN IF    PRI STATE          BW (kbits/sec)
------------------------ ---------- ---------- --- -------------- --------------
1.1.1.1 0_2              -          Gi0/1/0/0  1/1 Resv Admitted 50           RB0

RP/0/5/CPU0:R1#Sh rsvp int g0/1/0/0

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

Interface    MaxBW (bps)  MaxFlow (bps) Allocated (bps)      MaxSub (bps)
------------ ------------ ------------- -------------------- -------------
Gi0/1/0/0            80K             80K            50K ( 62%)             0*
```

30K more can be allocated on this interface.

Task 2

Configure a second tunnel called Tunnel "10 "to use 40 K of bandwidth. The priority should be identical to tunnel 0. Traffic destined for 44.4.4.4 should use this tunnel instead of tunnel "0". Policy-Based routing is not permitted.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On R1

```
RP/0/5/CPU0:R1(config)#int tunnel-te10
RP/0/5/CPU0:R1(config-if)#ipv4 unnumbered lo0
RP/0/5/CPU0:R1(config-if)#destination 4.4.4.4
RP/0/5/CPU0:R1(config-if)#priority 1 1
RP/0/5/CPU0:R1(config-if)#signalled-bandwidth 40
RP/0/5/CPU0:R1(config-if)#path-option 1 dynamic
RP/0/5/CPU0:R1(config-if)#commit
RP/0/5/CPU0:R1(config-if)#root

RP/0/5/CPU0:R1#Sh rsvp int

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

Interface   MaxBW (bps)  MaxFlow (bps) Allocated (bps)      MaxSub (bps)
----------- ------------ ------------- -------------------- ------------
Gi0/1/0/0           80K             80K           50K ( 62%)            0*
Gi0/1/0/2           80K             80K           40K ( 50%)            0*
Gi0/1/0/3          768K            768K             0 (  0%)            0*
```

The bandwidth constraint has been raised to 40 to force a calculation to be made on which tunnel gets to use the best path. The primary path has now been overutilized by 10K.

```
RP/0/5/CPU0:R1(config)#router static
RP/0/5/CPU0:R1(config-static)#address-family ipv4 unicast
RP/0/5/CPU0:R1(config-static-afi)#44.4.4.4/32 tunnel-te10
RP/0/5/CPU0:R1(config-static-afi)#commit
RP/0/5/CPU0:R1(config-static-afi)#end
```

This option can be used instead of routing traffic with autoroute announce

Verify Path Info:

```
RP/0/5/CPU0:R1#Sh cef exact-route 1.1.1.1 44.4.4.4

44.4.4.4/32, version 276, attached, internal 0x4000081 (ptr 0xae186980)
[1], 0x0 (0xae154f50), 0x0 (0x0)
 Updated May 16 02:46:05.842
 Prefix Len 32, traffic index 0, precedence routine (0), priority 2
   via tunnel-te10
   via tunnel-te10, 3 dependencies, weight 0, class 0 [flags 0x8]
    path-idx 0 [0xaea21280 0x0]
    local adjacency
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R1#Sh cef exact-route 1.1.1.1 4.4.4.4

4.4.4.4/32, version 270, internal 0x4000001 (ptr 0xae186728) [1], 0x0
(0xae154c08), 0x0 (0x0)
 Updated May 16 02:38:46.422
 Prefix Len 32, traffic index 0, precedence routine (0), priority 1
   via tunnel-te0
   via 4.4.4.4, tunnel-te0, 3 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 [0xaea211c8 0x0]
    next hop 4.4.4.4
    local adjacency
```

The "autoroute announce" option on "Tunnel 0" allows ALL traffic that is destined for 4.4.4.4 to use tunnel 0. The static overrides this constraint by choosing to use "Tunnel 10" for routes to the 44.4.4.4 destination prefix.

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels name tunnel-te10 | b Path info

  Path info (OSPF 1 area 0):
  Node hop count: 2
  Hop0: 10.1.12.2
  Hop1: 10.1.24.2
  Hop2: 10.1.24.4
  Hop3: 4.4.4.4
Displayed 1 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads

RP/0/5/CPU0:R1#Sh rsvp int

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

Interface   MaxBW (bps)  MaxFlow (bps) Allocated (bps)      MaxSub (bps)
----------- ------------ ------------- -------------------- ------------
Gi0/1/0/0           80K            80K           50K ( 62%)            0*
Gi0/1/0/2           80K            80K           40K ( 50%)            0*
Gi0/1/0/3          768K           768K             0 (  0%)            0*
```

The path through R3 was chosen due to the fact that Tunnel 0 was already established with a tunnel priority of "1" for setup and "1" for hold. "Tunnel 10" attempts to signal but the bandwidth constraint are not sufficient for the path to R4 directly. The tunnel reroutes all traffic through R2 since it is the next best path with the most available.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

Task 3

Tunnel 10 should take the path through R1, R3 and R4. Priority, Policy based Routing explicit paths or modifying OSPF cost is not permitted.

Viewing the current tunnel state:

```
RP/0/5/CPU0:R1#Sh mpls traffic-eng link-management advertisements | b Link ID:: 1

      Link ID:: 1 (GigabitEthernet0/1/0/2)
          Link IP Address      : 10.1.12.1
          O/G Intf ID          : 11
          Designated Router    : 10.1.12.1
          TE Metric            : 64
          IGP Metric           : 64
          Physical BW          : 1000000 kbits/sec
          BCID                 : RDM
          Max Reservable BW    : 80 kbits/sec
          Res Global BW        : 80 kbits/sec
          Res Sub BW           : 0 kbits/sec

RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels name tunnel-te10 | i Metric Type

    Metric Type: TE (default)
```

TE metric is another cost that can be used instead of the IGP metric. By default, the tunnel TE metrics also equals the IGP metric. (TE = IGP).

On R3

```
R3(config)#interface g0/0
R3(config-subif)#mpls traffic-eng administrative-weight 1
```

On R1

```
RP/0/5/CPU0:R1(config)#mpls traffic-eng
RP/0/5/CPU0:R1(config-mpls-te)#int g0/1/0/3
RP/0/5/CPU0:R1(config-mpls-te-if)#admin-weight 1
RP/0/5/CPU0:R1(config-mpls-te-if)#commit
RP/0/5/CPU0:R1(config-mpls-te-if)#root

RP/0/5/CPU0:R1(config)#int tunnel-te0
RP/0/5/CPU0:R1(config-if)#path-selection metric igp
RP/0/5/CPU0:R1(config-if)#commit
RP/0/5/CPU0:R1(config-if)#end

RP/0/5/CPU0:R1(config)#int tunnel-te0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:R1(config-if)#path-selection metric te
RP/0/5/CPU0:R1(config-if)#commit
RP/0/5/CPU0:R1(config-if)#end

R1#mpls traffic-eng reoptimize

RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels name tunnel-te0 | b Path info

  Path info (OSPF 1 area 0):
  Node hop count: 1
  Hop0: 10.1.14.4
  Hop1: 4.4.4.4
Displayed 1 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads

RP/0/5/CPU0:R1#Sh mpls traffic-eng tunnels name tunnel-te10 | b Path info

  Path info (OSPF 1 area 0):
  Node hop count: 2
  Hop0: 10.1.13.3
  Hop1: 10.1.34.3
  Hop2: 10.1.34.4
  Hop3: 4.4.4.4
Displayed 1 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
Displayed 1 up, 0 down, 0 recovering, 0 recovered heads

R1#Show mpls traffic-eng link-managemen interfaces s0/0.13 | incl Admin.
Weight:

Admin. Weight:          1 (configured)

RP/0/5/CPU0:R1#Sh mpls traffic-eng link-management interfaces g0/1/0/3 | i Admin

            Admin weight: 1 (TE), 130 (IGP)
```

Tunnel "0" uses the path through the lowest IGP cost which still uses the path to R4 directly.
Tunnel"10" uses the path that reflects the lowest TE metric of "2".

Task 4

Stop all routers in the console window and exit. Stop the Server and press the "Erase Start" launcher before proceeding.

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

Advanced
CCIE SERVICE PROVIDER
v3.0

www.MicronicsTraining.com

Narbik Kocharians
CCIE #12410
R&S, Security, SP

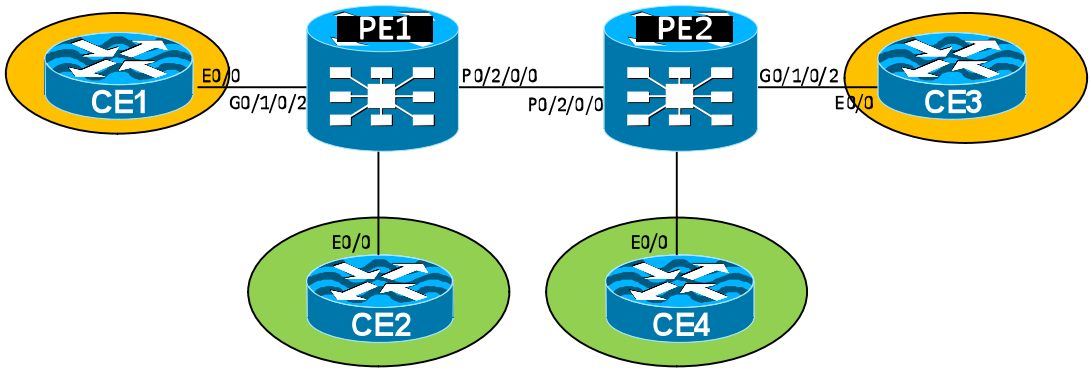
Paul Negron
CCIE #14856
SP

VPWS & VPLS

©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

LAB 1- VPWS



Router	Interface	IP address
CE1	Lo0 E0/0	11.11.11.11 /32 10.13.13.1/ 24
CE2	Lo0 E0/0	22.22.22.22/ 32 10.24.24.2/ 24
CE3	Lo0 E0/0	33.33.33.33 /32 10.13.13.3/ 24
CE4	Lo0 E0/0	44.44.44.44/ 32 10.24.24.4/ 24
PE1	Lo0 POS0/2/0/0	1.1.1.1/32 10.1.12.1/30
PE2	Lo0 POS0/2/0/0	2.2.2.2/32 10.1.12.2/30

Lab Setup:

- Configure IP connectivity according to the table above.

Task 1

Prepare the PE routers to support Layer 2 services on the interfaces closest to CE routers. These interfaces should be configured as an access port for vlan 10.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On PE-1

```
RP/0/5/CPU0:PE-1(config)#int g0/1/0/2.10 l2transport
RP/0/5/CPU0:PE-1(config-subif)#dot1q vlan 10
RP/0/5/CPU0:PE-1(config-subif)#commit
```

On PE-2

```
RP/0/5/CPU0:PE-2(config)#int g0/1/0/2.10 l2transport
RP/0/5/CPU0:PE-2(config-subif)#dot1q vlan 10
RP/0/5/CPU0:PE-2(config-subif)#commit
```

The interfaces for IOS-XR are configured with the l2transport to signify use. IOS would have you configure "encapsulation dot1q <vlan-id>" prior to configuring an ip address fro layer 2 use.

Task 2

Configure ISIS as the routing protocol between PE-1 and PE-2 for ip reachabiltiy between loopacks.

On PE-1

```
RP/0/5/CPU0:PE-1(config)#router isis TST
RP/0/5/CPU0:PE-1(config-isis)#net 49.0012.1111.0000.1111.00

RP/0/5/CPU0:PE-1(config-isis)#address-family ipv4 unicast
RP/0/5/CPU0:PE-1(config-isis-af)#metric-style wide

RP/0/5/CPU0:PE1(config-isis-af)#int POS0/2/0/0
RP/0/5/CPU0:PE-2(config-isis-if)#address-family ipv4 uni

RP/0/5/CPU0:PE1(config-isis-if-af)#int lo0
RP/0/5/CPU0:PE-1(config-isis-if)#passive
RP/0/5/CPU0:PE-2(config-isis-if)#address-family ipv4 uni

RP/0/5/CPU0:PE-1(config-isis-if)#commit
```

On PE-2

```
RP/0/5/CPU0:PE-2(config)#router isis TST
RP/0/5/CPU0:PE-2(config-isis)#net 49.0012.2222.0000.2222.00
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:PE-2(config-isis)#address-family ipv4 unicast
RP/0/5/CPU0:PE-2(config-isis-af)#metric-style wide

RP/0/5/CPU0:PE-2(config-isis-af)#int POS0/2/0/0
RP/0/5/CPU0:PE-2(config-isis-if)#address-family ipv4 uni

RP/0/5/CPU0:PE-2(config-isis-if-af)#int lo0
RP/0/5/CPU0:PE-2(config-isis-if)#passive
RP/0/5/CPU0:PE-2(config-isis-if)#address-family ipv4 uni

RP/0/5/CPU0:PE-2(config-isis-if-af)#commit
```

ISIS is configured with the "address-family unicast" option under the process to start it as well as each interface.

Verifying Connectivity:

On PE-1

```
RP/0/5/CPU0:PE1#Sh isis neighbors

IS-IS TST neighbors:
System Id      Interface        SNPA           State Holdtime Type IETF-NSF
PE2            PO0/2/0/0        *PtoP*         Up    22       L1L2 Capable

RP/0/5/CPU0:PE-1#Sh route isis

i L1 2.2.2.2/32 [115/10] via 10.1.12.2, 00:01:20, POS0/2/0/0

RP/0/5/CPU0:PE-1#Ping 2.2.2.2 so lo0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2.2.2.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 2/3/9 ms
```

On CE1

```
PE1-SW#Ping 10.13.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.13.13.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

The ping is verified to fail from CE1 to CE3. A pseudowire service is needed in the backbone for

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

connectivity between CE's.

Task 3

Create connectivity between the the VLAN 10 interafaces configured on the Switches across Backbone Routers PE-1 and PE-2 using a Wire Service that is not bridged.

On PE-1

```
RP/0/5/CPU0:PE-1(config)#l2vpn

RP/0/5/CPU0:PE-1(config-l2vpn)#pw-class TST-PE1
RP/0/5/CPU0:PE-1(config-l2vpn-pwc)#encapsulation mpls
RP/0/5/CPU0:PE-1(config-l2vpn-pwc-mpls)#l2vpn

RP/0/5/CPU0:PE-1(config-l2vpn)#xconnect group TST-GROUP
RP/0/5/CPU0:PE-1(config-l2vpn-xc)#p2p to-PE2
RP/0/5/CPU0:PE-1(config-l2vpn-xc-p2p)#int g0/1/0/2.10
RP/0/5/CPU0:PE-1(config-l2vpn-xc-p2p)#neighbor 2.2.2.2 pw-id 1002 pw-class TST-PE1
RP/0/5/CPU0:PE-1(config-l2vpn-xc-p2p)#commit
RP/0/5/CPU0:PE-1(config-l2vpn-xc-p2p)#end
```

L2vpn is a specific mode in IOS-XR. It allows the user to take advantage of the pseudowire service.

The pw-class is configured in a similar fashion to that of IOS.

The "xconnect group" allows us to take advantage of configuring a number of pseudowires that are grouped together for a purpose.

```
RP/0/5/CPU0:PE-1#Sh l2vpn xconnect | begin X

XConnect                    Segment 1                    Segment 2
Group      Name       ST    Description            ST    Description            ST
------------------------    --------------------------   --------------------------
TST-GROUP  to-PE2     DN    Gi0/1/0/2.10           UP    2.2.2.2          1002  DN

RP/0/5/CPU0:PE-1#Sh l2vpn xconnect group TST-GROUP detail

Group TST-GROUP, XC to-PE2, state is down; Interworking none
  AC: GigabitEthernet0/1/0/2.10, state is up
    Type VLAN; Num Ranges: 1
    VLAN ranges: [10, 10]
    MTU 1500; XC ID 0x2000004; interworking none
    Statistics:
      packets: received 0, sent 0
      bytes: received 0, sent 0
```


©2012 Micronics Networking & Training Inc. All rights reserved

Surinderjeet Singhh

```
      drops: illegal VLAN 0, illegal length 0
  PW: neighbor 2.2.2.2, PW ID 1002, state is down ( provisioned )
    PW class TST-PE1, XC ID 0x2000004
    Encapsulation MPLS, protocol LDP
    PW type Ethernet VLAN, control word disabled, interworking none
    PW backup disable delay 0 sec
    Sequencing not set

      MPLS         Local                          Remote
      ------------ ------------------------------ ------------------------
      Label        16000                          unknown
      Group ID     0x2000500                      0x0
      Interface    GigabitEthernet0/1/0/2.10      unknown
      MTU          1500                           unknown
```

The Attached Circuit is UP although the path is down. RSVP or

```
RP/0/5/CPU0:PE-1#Sh l2vpn xconnect brief

AToM
  Like-to-Like                    UP       DOWN        UNR
    VLAN                           0          1          0
    Total                          0          1          0

  Total                            0          1          0
```

On PE-2

```
RP/0/5/CPU0:PE-2(config-l2vpn)#l2vpn
RP/0/5/CPU0:PE-2(config-l2vpn)#pw-class TST-PE2
RP/0/5/CPU0:PE-2(config-l2vpn-pwc)#encap mpls
RP/0/5/CPU0:PE-2(config-l2vpn-pwc-mpls)#l2vpn
RP/0/5/CPU0:PE-2(config-l2vpn)#xconnect group TST-GROUP
RP/0/5/CPU0:PE-2(config-l2vpn-xc)#p2p to-PE1
RP/0/5/CPU0:PE-2(config-l2vpn-xc-p2p)#int g0/1/0/2.10

RP/0/5/CPU0:PE-2(config-l2vpn-xc-p2p)#neighbor 1.1.1.1 pw-id 1002 pw-class TST-PE2
RP/0/5/CPU0:PE-2(config-l2vpn-xc-p2p)#commit

RP/0/5/CPU0:PE-2#Sh l2vpn xconnect group TST-GROUP detail

Group TST-GROUP, XC to-PE1, state is down; Interworking none
  AC: GigabitEthernet0/1/0/2.10, state is up
    Type VLAN; Num Ranges: 1
    VLAN ranges: [10, 10]
    MTU 1500; XC ID 0x2000002; interworking none
    Statistics:
      packets: received 0, sent 0
      bytes: received 0, sent 0
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
      drops: illegal VLAN 0, illegal length 0
  PW: neighbor 1.1.1.1, PW ID 1002, state is down ( provisioned )
    PW class TST-PE2, XC ID 0x2000002
    Encapsulation MPLS, protocol LDP
    PW type Ethernet VLAN, control word disabled, interworking none
    PW backup disable delay 0 sec
    Sequencing not set

RP/0/5/CPU0:PE-2#Sh l2vpn xconnect group TST-GROUP | b X

XConnect                    Segment 1                   Segment 2
Group      Name       ST    Description            ST   Description             ST
-----------------------    -------------------------   -------------------------
TST-GROUP  to-PE1     DN    Gi0/1/0/2.10           UP   1.1.1.1         1002    DN
```

On PE-1

```
RP/0/5/CPU0:PE-1(config)#mpls ldp
RP/0/5/CPU0:PE-1(config-ldp)#router-id 1.1.1.1
RP/0/5/CPU0:PE-1(config-ldp)#int pos0/2/0/0
RP/0/5/CPU0:PE-1(config-ldp-if)#commit
```

On PE-2

```
RP/0/5/CPU0:PE-2(config)#mpls ldp
RP/0/5/CPU0:PE-2(config-ldp)#router-id 2.2.2.2
RP/0/5/CPU0:PE-2(config-ldp)#int pos0/2/0/0
RP/0/5/CPU0:PE-2(config-ldp-if)#commit
```

On PE-2

```
RP/0/5/CPU0:PE-2#Sh l2vpn xconnect group TST-GROUP | b X

XConnect                    Segment 1                   Segment 2
Group      Name       ST    Description            ST   Description             ST
-----------------------    -------------------------   -------------------------
TST-GROUP  to-PE1     UP    Gi0/1/0/2.10           UP   1.1.1.1         1002    UP
```

On PE-1

```
RP/0/5/CPU0:PE-1#Sh l2vpn xconnect group TST-GROUP detail

Group TST-GROUP, XC to-PE2, state is up; Interworking none
  AC: GigabitEthernet0/1/0/2.10, state is up
    Type VLAN; Num Ranges: 1
    VLAN ranges: [10, 10]
    MTU 1500; XC ID 0x2000004; interworking none
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    Statistics:
      packets: received 0, sent 3
      bytes: received 0, sent 204
      drops: illegal VLAN 0, illegal length 0
  PW: neighbor 2.2.2.2, PW ID 1002, state is up ( established )
    PW class TST-PE1, XC ID 0x2000004
    Encapsulation MPLS, protocol LDP
    PW type Ethernet VLAN, control word disabled, interworking none
    PW backup disable delay 0 sec
    Sequencing not set

      MPLS         Local                          Remote
      ------------ ------------------------------ ------------------------
      Label        16000                          16000
      Group ID     0x2000500                      0x2000500
      Interface    GigabitEthernet0/1/0/2.10
GigabitEthernet0/1/0/2.10
      MTU          1500                           1500
      Control word disabled                       disabled
      PW type      Ethernet VLAN                  Ethernet VLAN
      VCCV CV type 0x2                            0x2
                   (LSP ping verification)        (LSP ping verification)
      VCCV CC type 0x6                            0x6
                   (router alert label)           (router alert label)
                   (TTL expiry)                   (TTL expiry)
      ------------ ------------------------------ ------------------------
    MIB cpwVcIndex: 3
    Create time: 19/03/2012 20:05:51 (00:13:31 ago)
    Last time status changed: 19/03/2012 20:19:09 (00:00:12 ago)
    Statistics:
      packets: received 3, sent 0
      bytes: received 204, sent 0
```

On CE1

```
CE1#Ping 10.13.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.12.12.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms
```

On CE2

```
CE3#Ping 10.24.24.4
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.34.34.4, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

Task 4

Create connectivity between two VLAN 20 interfaces configured on Switches PE3-SW and PE3-SW as you did in the prior task. Connectivity between CE3 and CE 4 should be established using ospf as a routing protocol.

On PE-1

```
RP/0/5/CPU0:PE-1(config)#int g0/1/0/2.20 l2transport
RP/0/5/CPU0:PE-1(config-subif)#dot1q vlan 20
RP/0/5/CPU0:PE-1(config)#root
RP/0/5/CPU0:PE-1(config)#l2vpn
RP/0/5/CPU0:PE-1(config-l2vpn)#xconnect group TST-GROUP
RP/0/5/CPU0:PE-1(config-l2vpn-xc)#p2p to-PE2-2
RP/0/5/CPU0:PE-1(config-l2vpn-xc-p2p)#int g0/1/0/2.20
RP/0/5/CPU0:PE-1(config-l2vpn-xc-p2p)#neighbor 2.2.2.2 pw-id 1022 pw-class TST-PE1
RP/0/5/CPU0:PE-1(config-l2vpn-xc-p2p)#commit
```

On PE-2

```
RP/0/5/CPU0:PE-2(config)#int g0/1/0/2.20 l2transport
RP/0/5/CPU0:PE-2(config-subif)#dot1q vlan 20
RP/0/5/CPU0:PE-2(config)#root

RP/0/5/CPU0:PE-2(config)#l2vpn
RP/0/5/CPU0:PE-2(config-l2vpn)#xconnect group TST-GROUP
RP/0/5/CPU0:PE-2(config-l2vpn-xc)#p2p to-PE1-2
RP/0/5/CPU0:PE-2(config-l2vpn-xc-p2p)#int g0/1/0/2.20
RP/0/5/CPU0:PE-2(config-l2vpn-xc-p2p)#neighbor 1.1.1.1 pw-id 1022 pw-class TST-PE2
RP/0/5/CPU0:PE-2(config-l2vpn-xc-p2p)#commit

RP/0/5/CPU0:PE-2#Sh l2vpn xconnect group TST-GROUP | b X

XConnect                    Segment 1                    Segment 2
Group      Name       ST    Description            ST    Description            ST
------------------------    --------------------------   --------------------------
TST-GROUP  to-PE1     UP    Gi0/1/0/2.10           UP    1.1.1.1          1002  UP

------------------------------------------------------------------------------------
TST-GROUP  to-PE1-2   UP    Gi0/1/0/2.20           UP    1.1.1.1          1022  UP
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On CE-2

```
PE1-SW#Ping 10.24.24.4

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.24.24.4, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/4/9 ms
```

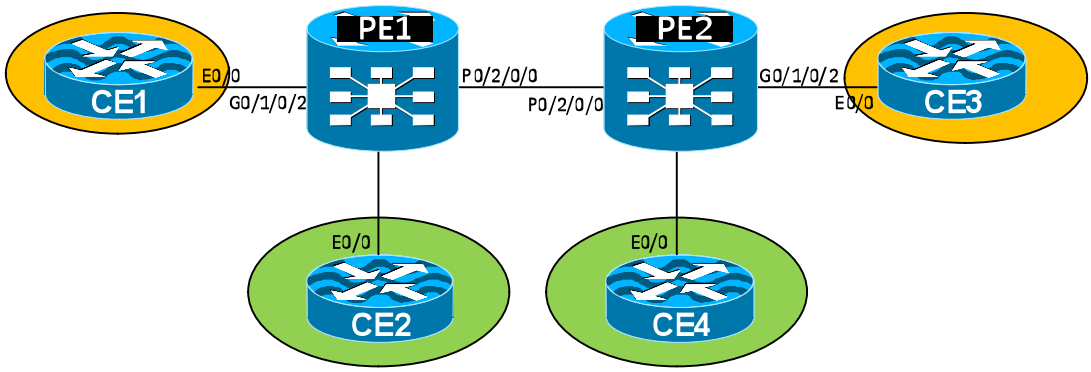
On CE-4

```
PE3-SW#Ping 10.24.24.2

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.24.24.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/4/9 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

LAB 2- VPLS (EP-LAN)



Router	Interface	IP address
CE1	Lo0 E0/0	11.11.11.11 /32 10.13.13.1/ 24
CE2	Lo0 E0/0	22.22.22.22/ 32 10.24.24.2/ 24
CE3	Lo0 E0/0	33.33.33.33 /32 10.13.13.3/ 24
CE4	Lo0 E0/0	44.44.44.44/ 32 10.24.24.4/ 24
PE1	Lo0 G0/1/0/0	1.1.1.1/32 10.1.12.1/30
PE2	Lo0 G0/1/0/0	2.2.2.2/32 10.1.12.2/30

Task 1

Prepare the Interface on PE1 and PE2 that connect to the CE routers to support Layer 2 connectivity for MPLS. Use VLAN 10 for the PVC between CE1 and CE3.

On PE1

```
RP/0/5/CPU0:PE1(config)#int g0/1/0/2
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:PE1(config-if)#No shut
RP/0/5/CPU0:PE1(config-if)#int g0/1/0/2.10 l2transport
RP/0/5/CPU0:PE1(config-subif)#dot1q vlan 10
RP/0/5/CPU0:PE1(config-subif)#commit
RP/0/5/CPU0:PE1(config-subif)#end
```

On PE2

```
RP/0/5/CPU0:PE2(config)#int g0/1/0/2
RP/0/5/CPU0:PE2(config-if)#No shut
RP/0/5/CPU0:PE2(config-if)#int g0/1/0/2.10 l2transport
RP/0/5/CPU0:PE2(config-subif)#dot1q vlan 10
RP/0/5/CPU0:PE2(config-subif)#commit
RP/0/5/CPU0:PE2(config-subif)#end
RP/0/5/CPU0:PE-2(config-subif)#commit
```

There is no concept of as a vlan interface on the IOS-XR platform. Instead, a subinterface type is declared after the logical name "interface g0/1/0/2.10".

Task 2

Configure ISIS as the routing protocol between PE1 and PE-2 for ip reachabiltiy between loopbacks.

On PE1

```
RP/0/5/CPU0:PE1(config)#router isis TST
RP/0/5/CPU0:PE1(config-isis)#net 49.0012.1111.0000.1111.00

RP/0/5/CPU0:PE1(config-isis)#address-family ipv4 unicast
RP/0/5/CPU0:PE1(config-isis-af)#metric-style wide

RP/0/5/CPU0:PE1(config-isis-af)#int G0/1/0/0
RP/0/5/CPU0:PE-2(config-isis-if)#address-family ipv4 uni

RP/0/5/CPU0:PE1(config-isis-if)#int lo0
RP/0/5/CPU0:PE1(config-isis-if)#passive
RP/0/5/CPU0:PE-2(config-isis-if)#address-family ipv4 uni

RP/0/5/CPU0:PE1(config-isis-if)#commit
```

On PE-2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:PE-2(config)#router isis TST
RP/0/5/CPU0:PE-2(config-isis)#net 49.0012.2222.0000.2222.00

RP/0/5/CPU0:PE-2(config-isis)#address-family ipv4 unicast
RP/0/5/CPU0:PE-2(config-isis-af)#metric-style wide

RP/0/5/CPU0:PE-2(config-isis-af)#int g0/1/0/0
RP/0/5/CPU0:PE-2(config-isis-if)#address-family ipv4 uni

RP/0/5/CPU0:PE-2(config-isis-if-af)#int lo0
RP/0/5/CPU0:PE-2(config-isis-if)#passive
RP/0/5/CPU0:PE-2(config-isis-if)#address-family ipv4 uni

RP/0/5/CPU0:PE-2(config-isis-if-af)#commit
```

Verifying Connectivity:

On PE2

```
RP/0/5/CPU0:PE2#Sh route isis

i L1 1.1.1.1/32 [115/10] via 10.1.12.1, 00:00:43, POS0/2/0/0

RP/0/5/CPU0:PE2#Ping 1.1.1.1 so loopback0

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/44/214 ms
```

Task 3

Create connectivity between the the VLAN 10 interaces configured on the Switches across Backbone Routers PE1 and PE-2 using a LAN Service that is BRIDGED.

On PE1

```
RP/0/5/CPU0:PE1(config)#l2vpn

RP/0/5/CPU0:PE1(config-l2vpn)#pw-class TST-PE1
RP/0/5/CPU0:PE1(config-l2vpn-pwc)#encapsulation mpls
RP/0/5/CPU0:PE1(config-l2vpn-pwc-mpls)#exit
RP/0/5/CPU0:PE1(config-l2vpn-pwc)#exit
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

The pw-class is configured to note an MPLS encapsulation.

```
RP/0/5/CPU0:PE1(config-l2vpn)#bridge group PE1-PE2
RP/0/5/CPU0:PE1(config-l2vpn-bg)#bridge-domain VLAN10
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd)#interface G0/1/0/2.10
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd-ac)#exit
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd)#VFI TSTA
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd-vfi)#neighbor 2.2.2.2 pw-id 1002 pw-class TST-PE1
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd-vfi)#commit
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd-vfi)#end
```

A bridge group must be created to contain a number of bridge domains to assign interfaces to. A Virtual Forwarding Interface (VFI) receives frames from interfaces assigned to a bridge domain and forwards them in a unidirectional manner to another VFI peer. Traffic received from a VFI peer is not permitted to be forwarded to another VFI peer by default. (Split-Horizon)

On PE-2

```
RP/0/5/CPU0:PE-2(config)#l2vpn

RP/0/5/CPU0:PE-2(config-l2vpn)#pw-class TST-PE2
RP/0/5/CPU0:PE-2(config-l2vpn-pwc)#encapsulation mpls
RP/0/5/CPU0:PE-2(config-l2vpn-pwc-mpls)#l2vpn

RP/0/5/CPU0:PE2(config-l2vpn)#bridge group PE1-PE2
RP/0/5/CPU0:PE2(config-l2vpn-bg)#bridge-domain VLAN10
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd)#interface G0/1/0/2.10
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd-ac)#exit
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd)#VFI TSTA
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd-vfi)#neighbor 1.1.1.1 pw-id 1002 pw-class TST-PE2
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd-vfi)#commit
```

On PE1

```
RP/0/5/CPU0:PE1#Sh l2vpn bridge-domain brief

Bridge Group/Bridge-Domain Name  ID    State      Num ACs/up     Num PWs/up
-------------------------------- ----- ---------- -------------- --------------
PE1-PE2/VLAN10                   0     up         1/0            1/0
```

On PE1

```
RP/0/5/CPU0:PE1(config)#mpls ldp
RP/0/5/CPU0:PE1(config-ldp)#router-id 1.1.1.1
RP/0/5/CPU0:PE1(config-ldp)#int pos0/2/0/0
RP/0/5/CPU0:PE1(config-ldp-if)#commit
```

On PE2

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
RP/0/5/CPU0:PE-2(config)#mpls ldp
RP/0/5/CPU0:PE-2(config-ldp)#router-id 2.2.2.2
RP/0/5/CPU0:PE-2(config-ldp)#int pos0/2/0/0
RP/0/5/CPU0:PE-2(config-ldp-if)#commit
```

After configuring the path of the MPLS circuit the Layer 2 VPN is complete.

On PE1

```
RP/0/5/CPU0:PE1#Sh l2vpn bridge-domain brief

Bridge Group/Bridge-Domain Name  ID    State      Num ACs/up     Num PWs/up
-------------------------------- ----- ---------- -------------- --------------
PE1-PE2/VLAN10                   0     up         1/1            1/1
```

On PE2

```
RP/0/5/CPU0:PE2#Sh l2vpn bridge-domain bd-name VLAN10

Bridge group: PE2-PE1, bridge-domain: VLAN10, id: 0, state: up, ShgId: 0, MSTi: 0
  Aging: 300 s, MAC limit: 4000, Action: none, Notification: syslog
  Filter MAC addresses: 0
  ACs: 1 (1 up), VFIs: 1, PWs: 1 (1 up), PBBs: 0 (0 up)
  List of ACs:
    Gi0/1/0/2.10, state: up, Static MAC addresses: 0
  List of Access PWs:
  List of VFIs:
    VFI TSTA
      Neighbor 1.1.1.1 pw-id 1002, state: up, Static MAC addresses: 0
```

On CE1

```
CE1#Ping 10.13.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.13.13.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/3/9 ms
```

Task 4

Create connectivity between the the VLAN 20 interfaces configured on Switches PE1 and PE2 as you did in the prior task.

©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

On PE1

```
RP/0/5/CPU0:PE1(config-l2vpn)#bridge group PE1-PE2-2
RP/0/5/CPU0:PE1(config-l2vpn-bg)#bridge-domain VLAN20
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd)#interface GigabitEthernet0/1/0/2.20
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd-ac)#vfi TSTB
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd-vfi)#neighbor 2.2.2.2 pw-id 1022
RP/0/5/CPU0:PE1(config-l2vpn-bg-bd-vfi-pw)#pw-class TST-PE1
```

On PE2

```
RP/0/5/CPU0:PE2(config)#l2vpn
RP/0/5/CPU0:PE2(config-l2vpn)#bridge group PE2-PE1-2
RP/0/5/CPU0:PE2(config-l2vpn-bg)#bridge-domain VLAN20
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd)#interface GigabitEthernet0/1/0/2.20
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd-ac)#vfi TSTB
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd-vfi)#neighbor 1.1.1.1 pw-id 1022
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd-vfi-pw)#pw-class TST-PE2
RP/0/5/CPU0:PE2(config-l2vpn-bg-bd-vfi-pw)#commit

RP/0/5/CPU0:PE2#Sh l2vpn bridge-domain bd-name VLAN20 brief
Bridge Group/Bridge-Domain Name  ID    State      Num ACs/up     Num PWs/up
-------------------------------- ----- ---------- -------------- --------------
PE2-PE1-2/VLAN20                 1     up         1/1            1/1
```

On CE3

```
CE3#Ping 10.34.34.4

Sending 5, 100-byte ICMP Echos to 10.20.20.2, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/2/8
```

On PE2

```
RP/0/5/CPU0:PE2#Sh l2vpn bridge-domain bd-name VLAN10

Tue Mar 20 01:30:42.393 UTC
Bridge group: PE2-PE1, bridge-domain: VLAN10, id: 0, state: up, ShgId: 0, MSTi: 0
  Aging: 300 s, MAC limit: 4000, Action: none, Notification: syslog
  Filter MAC addresses: 0
  ACs: 1 (1 up), VFIs: 1, PWs: 1 (1 up), PBBs: 0 (0 up)
  List of ACs:
    Gi0/1/0/2.10, state: up, Static MAC addresses: 0
  List of Access PWs:
  List of VFIs:
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh

```
    VFI TSTA
      Neighbor 1.1.1.1 pw-id 1002, state: up, Static MAC addresses: 0
```

On CE1

```
CE1#Ping 10.13.13.3

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.13.13.3, timeout is 2 seconds:
.!!!!
Success rate is 80 percent (4/5), round-trip min/avg/max = 1/3/9 ms
```


©2012 Micronics Networking & Training Inc. All rights reserved
Surinderjeet Singhh