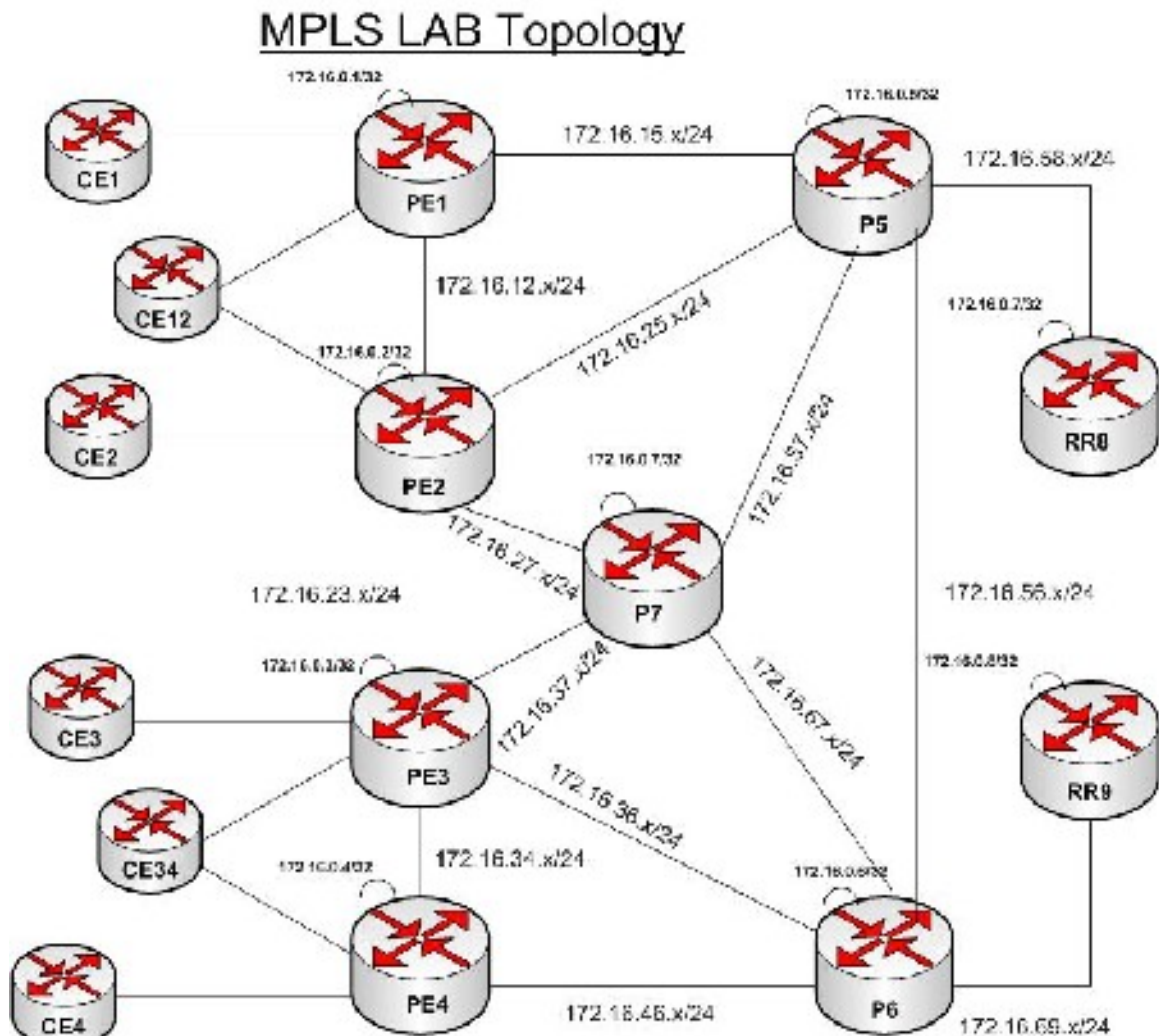


Lab 1:

Configuring MPLS

Configuring LDP and its optional parameters

Lab



Topology:

Objective 1: Network Familiarization

1. Familiarize yourself with the routers assigned to you.
2. Check the IGP Routes
3. Make sure you understand the network path between PE's

4. Telnet to other devices as well
5. Understand the IP addressing Scheme

Objective 2: Configuring LDP

6. Follow the steps from the presentation to enable LDP on your routers
7. Ensure that LDP session exists with your peering PE or P router
8. Understand the router who initiated the TCP session
9. Run debugs to check the discovery packets, what debugs did you use ?
10. Verify the protocol type (TCP or UDP) for LDP Discovery packets and the ports number
11. Verify the protocol type for LDP session establishment packets
12. Make sure you understand the network path between PE's
13. Trace the path between PE's , Are the packet forwarded as IP or Labeled packets
14. Check the Routing Table
15. Check the FIB Table
16. Check the LFIB Table
17. Check the LIB Table
18. Add a new loopback interface , Loopback 100 with ip address 100.100.100.x , where x is your Router number
19. Are the 100.100.100.x prefix available in the Routing table
20. Is the 100.100.100.x available in the LIB
21. Is the 100.100.100.x available in the LFIB
22. Does 100.100.100.x available in show mpls ldp neighbor detail
23. Change the LDP router id to the ip address configured on Loopback 100
24. Does that impact your neighbor relationship
25. Why cant we form LDP neighbor relationship
26. Perform the necessary steps to form the LDP neighbor relation ship with out adding static routes
27. Run different debugs and familiarize with the outputs
28. Enable **debug mpls ldp bindings on your router** , Shut down the interface Loopback 100 on your Router and check the debug logs
29. No shut the interface, do you see any debugs?
30. Change the mpls label protocol to TDP on the interface, do you now see LDP neighbor relation ship now?
31. Configure your interface to support both LDP and TDP relationship, Does that help?
32. Is your neighbor established via LDP or TDP?
33. Do you see ldp and TDP discovery packets?
34. Do you see any special tags (pop tags , Untagged etc) in the forwarding table
35. Configure your router to advertise an explicit null instead of implicit null to its LDP peers
36. What is the router ID of OSPF ? What command did you use
37. What is the router ID of LDP ? What command did you use

38. Modify your router id to use a physical interface address ? Use 'mpls ldp router-id fast 1/0 ' specify any working interface in 172.16.x.x subnet
39. Did the router id change right away?
40. Issue clear mpls ldp neighbor * and verify if the new router id is enforced
41. How can you change your router id ?
42. Configure Loopback 101 (do not configure any ip address)
43. Change your LDP router id to Loopback 101 Use 'mpls ldp router-id Loopback 101 force'
44. Did the router id change right away? If no why ?
45. Configure ip address 17.16.0.10x/32 on Loopback 101 , where x is your router number , for e.g for PE1 x=172.16.0.101/32
46. Does the mpls neighbors flap? Does the LDP router-id now change
47. Verify the LDP relationship is now up
48. Note the LDP router Id of the peering router
49. configure a static route of your peering router LDP ID and point it to null0
For e.g. if you peer with P6 then 'ip route 172.16.0.106 255.255.255.255 null 0'
50. Is the LDP neighbor relationship still up ?
51. clear mpls ldp neighbor * , Why didn't the neighbor disappear right away ?
52. Do you now see the neighbor via show mpls ldp neighbor
53. Issue show mpls ldp discovery , do you see the LDP discovery packets being sent and received ?
54. Why is the neighbor relationship not active?
55. Fix the problem by removing the route 'no ip route 172.16.0.106 255.255.255.255 null 0'
56. Change the router id of LDP back to Loopback 0 , make sure that the router id is now in effect.
57. Change the LDP hello interval to 1 sec , does that apply to all peers ? Use the command 'mpls ldp discovery hello interval 1'
58. Issue 'clear mpls ldp neighbor *' , check if the LDP hello interval of 1 sec is now negotiated ?
59. Change the LDP hello discovery hold time to 3 secs . Use the command 'mpls ldp discovery hello holdtime 3'
60. Does it change the hello interval
61. Issue 'clear mpls ldp neighbor *' , check if the LDP hello interval of 1 sec is now negotiated ?
62. Trace your path from PE to PE , Are the packets forwarded as IP or labeled packets?
63. Do you see all the routers in the path?
64. Configure the PE router so that it will not see any hop in between, What command did you use ?
65. Configure the subnet between CE's and PE's according to the following table
(Please use the subinterface with vlan id 1 and use that as your native vlan)

Add the description on PE and CE interfaces as
'description management of SP network'

CE name/Interface	IP address	PE name/Interface	IP address
CE1 (fa0/0.1)	192.168.11.11/24	PE1 (fa1/0.1)	192.168.11.1/24

CE2 (fa0/0.1)	192.168.12.12/24	PE2 (fa1/0.1)	192.168.12.2/24
CE12 (fa0/0.1)	192.168.112.112/24	PE1 (fa1/1.1)	192.168.112.1/24
CE12 (fa1/0.1)	192.168.212.112/24	PE2 (fa1/1.1)	192.168.212.2/24
CE3 (fa0/0.1)	192.168.33.13/24	PE3 (fa1/0.1)	192.168.33.3/24
CE4 (fa0/0.1)	192.168.44.14/24	PE4 (fa 1/0.1)	192.168.44.4/24
CE34 (fa 0/0.1)	192.168.3.34/24	PE3(fa 1/1.1)	192.168.3.3/24
CE34 (fa 1/0.1)	192.168.4.34/24	PE4 (fa 1/1.1)	192.168.3.4/24

66. ping the CE to local PE and ensure that the ip addresses are correctly given

67. Verify , your configuration looks as below

For e.g.

Rack2-CE12#show running-config interface fa 0/0.1

Building configuration...

Current configuration : 132 bytes

!

interface FastEthernet0/0.1

description management of SP network

encapsulation dot1Q 1 native

ip address 192.168.112.12 255.255.255.0

no snmp trap link-status

end

Rack2-PE1#show running-config interface fa 1/1.1

Building configuration...

Current configuration : 105 bytes

!

interface FastEthernet1/1.1

encapsulation dot1Q 1 native

ip address 192.168.112.1 255.255.255.0

end

68. Insert a static route on CE's and point to the PE (point to lower PE number if you have multiple nexthop PEs) for all 172.16.0.0/16 traffic

69. On the PE routers redistribute the connected routes into ospf . (Please ensure that you just redistribute the prefixes in 192.168.0.0/16 subnets)

70. The configuration on the PE will look as below

71. Rack2-PE1#show running-config interface fa 1/1.1

Rack2-PE1#show running-config interface fa 1/1.1

Building configuration...

Current configuration : 105 bytes

!

```
interface FastEthernet1/1.1
 encapsulation dot1Q 1 native
 ip address 192.168.112.1 255.255.255.0
end
```

72. Ensure that you see only the 192.16.x.x prefixes in ospf database , what command did you use
73. Issue show ip ospf database | be Type-5
74. On CE router. Ping loopback of PE's , P routers
75. On CE router, issue trace to PE's and P routers loopbacks , do you see the routers in the path. (You should only see 2 hops in most cases except for RR's)
76. From the PE , issue trace to other PE's , how many hops do you see
77. Configure the PE such that ping traffic from local PE sees all the MPLS hops where as the traffic generated by CE's only see 2 hops.

Objective : BGP Transit Scenario

1. Remove any loopbacks except for Loopback 0
2. Enable BGP Sessions between PEs and RR's (Use AS 100 , peer the sessions using Loopback0 as the update source , make sure you configure Loopback 0 as the BGP router-id)
3. Each PE should have 2 BGP sessions , one to each RR
4. Configure Loopback 100 on each PE with the ip addresses 100.100.x.1/24 , 100.101.x.1/24 , 100.102.x.1/24 , 100.103.x.2/24 , where x=PE#
5. Advertise these loopbacks into BGP (do not use any redistribute command)
6. Ensure that you see the loopback from other PE's in your Routing table
7. Ping from your Loopback 100 source ip address to remote PE's loopback 100 ip addresses , can you do that ?
8. Can you check the routing table on the P routers, do they have the route?
9. Are the ICMP packets generated on the PE sent out using local FIB or LFIB table?
10. Are the ICMP packets switched by the P routers using LFIB or FIB ?
11. Check the bytes counter on the MPLS LFIB entries on the P routers to verify the packets switched. Which specific LFIB entry shows the counters incrementing

Objective : Conditional Advertisement of Labels

1. Configure your PE and P routers such that they only send LDP bindings for the Loopbacks (in 172.16.0.x/24 subnet)
2. Do you need to advertise any LDP bindings from the P routers ?
3. Disable the LDP advertisement of any LDP bindings from the P routers , does it break any thing

4. Make sure end to end connectivity works (i.e you can ping the loopback 100 ip addresses between PE's) before you move to the next step

Objective: Traffic Black holing

1. Increase the ospf cost on the interfaces between P5-PE2 and P6-PE3
2. Make sure traffic from PE2 and PE3 to rest of PEs use the path via P7 , Issue show ip route <100.100.x.1> (remote PE loopback and notice the ospf cost)
3. Ensure the connectivity between PE's on Loopback 100 ip addresses is intact
4. Disable MPLS ip on all interfaces on P7
5. Check if the connectivity between PE's on Loopback 100 ip addresses is still intact
6. What is the root cause of the issue? Please explain with outputs captured
7. What can be the options to fix this (Do not fix the issue) . Just write the options down
8. Without enabling mpls ip on P7 and changing any ospf costs , think of an alternative
9. On all PE and P's, enable 'mpls ldp sync' feature in your ospf process
10. See if that helps to restore the connectivity
11. Remove mpls ldp sync feature from PE1 and PE4 and if you still any issues for the traffic between all PEs (To Loopback 100's communication)
12. Configure IGP LDP synchronization feature on all P and PE's

Objective: Reconfiguring MPLS on P7 and rest of PE's

1. Configure one command to enable LDP on all interfaces of P7
2. Use the same command on all P and PE routers
3. Verify the packets (on Loopback 100 between PEs) now chose path via P7

Objective: Using BGP as a Label distribution protocol

1. Issue show ip bgp 100.100.x.0 and show ip bgp 100.100.y.0 , where x is a remote PE and y is the local PE , capture the output in a notepad
2. Issue show mpls forwarding 100.100.x.0 det and show mpls forwarding 100.100.y.0 det,
3. Reconfigure the BGP connectivity on PEs and RRs to send labels via BGP.
4. Issue show ip bgp 100.100.x.0 and show ip bgp 100.100.y.0 , where x is a remote PE and y is the local PE , capture the output in a notepad
5. Issue show mpls forwarding 100.100.x.0 det and show mpls forwarding 100.100.y.0 det,
6. Do you see an implicit null label advertised via BGP
7. Wait for your peers to finish

8. Reconfigure the BGP connectivity on PEs and RRs to send explicit null label via BGP
9. Issue show ip bgp 100.100.x.0 and show ip bgp 100.100.y.0 , where x is a remote PE and y is the local PE , capture the output in a notepad
10. Issue show mpls forwarding 100.100.x.0 det and show mpls forwarding 100.100.y.0 det,
11. Capture show mpls interface detail
- 12.

Objective: Configuring targeted neighbor sessions

1. Configure LDP relationship between PE's
2. Verify the neighbors are up