MPLS Network Design and Deployment Workshop Information

Introduction

Prerequisites

Prospective participants of this workshop should have completed the Cisco Internetworking Courses or have experience configuring, maintaining, and troubleshooting Cisco devices. At a minimum, participants should have an understanding of, and some experience with, the following:

- Use of Access-lists, Route-maps
- Network design familiarity
- Familiarity with LAN/WAN technologies
- Cisco router maintenance basics
- Design and implementation of IP addressing schemes
- Routing Information Protocol (RIP) basics
- Interior Gateway Routing Protocol (IGRP) basics
- Enhanced Interior Gateway Routing Protocol (EIGRP) basics
- Open Shortest Path First (OSPF) basics
- Border Gateway Protocol (BGP) basics

Each of the modules in this course contains an intensively detailed hands-on live lab exercise. These exercises will reinforce the concepts and skills covered in the lectures and will offer you an opportunity to truly apply these concepts to live/emulated Cisco equipment. You will have the opportunity to directly apply configuration and troubleshooting skills in a very realistic network environment.

Workshop Structure

This workshop is divided into several distinct modules. It is highly recommended that you complete all of the exercises within each module before moving on to the next. This will help in immediately reinforcing lecture/presentation material with hands-on lab exercises, and will also help you progress in a systematic manner through the course. Also, since you can only have one live lab exercises running at any given time, it is suggested that you complete the lab (or at least save your work) before moving on to the next protocol.

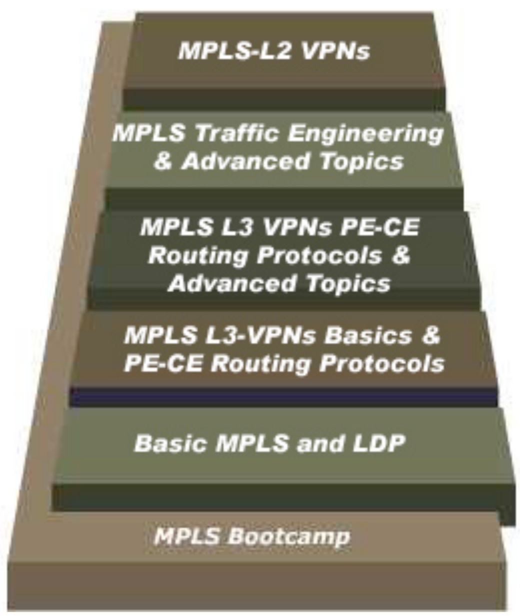

For each module, after listening to the lecture material, you should progress into the hands-on lab exercises, which give you a chance to truly apply and further your comprehension of the presented lecture material. So, mode for each module would be lecture material followed by hands-on lab in the below format:

Module #1: MPLS Basics and LDP

Overview

This module covers MPLS fundamentals including the following topics:

- Why MPLS is needed
- What protocols are used to distribute labels
- How labels are advertised and stored
- LDP, LDP concepts, configuration and troubleshooting

Objectives

After completing this module, you should be able to:

- Understand MPLS concepts and protocols.
- Understand the difference between: (1) Downstream on demand vs. Downstream unsolicited (2) Liberal vs. Conservative label retention (3) Independent vs. Ordered Control.
- Apply MPLS concepts to network design.
- Explain why you would use MPLS and troubleshoot MPLS scenarios.
- Understand, configure and troubleshoot LDP.

After completing this module, you should be able to:

- Understand label assignment and distribution.
- Describe Label Switched Paths (LSP).
- Explain loop detection and prevention.
- Characterize Label Distribution Protocol (LDP).
- Configure and monitor MPLS on Cisco IOS Frame-Mode interfaces.

LAB 1: Implementing Multiprotocol Label Switching (MPLS)
LAB 2: Establishing the Core MPLS Environment

LAB 1: Implementing Multiprotocol Label Switching (MPLS)

Introduction: Introduction

Welcome to MPLS Lab 1: Implementing Multiprotocol Label Switching (MPLS). Please progress through the steps in this lab, performing the configuration and troubleshooting exercises found in each step. A topology for the lab is supplied in order to give you a visual representation of the devices you are working with. Use the supplied interface to interact with the devices and complete the lab.

Step 1: Important Information

.

Router Naming Convention:
Service Provider P routers – P5, P6, P7 (Simulation R5, R6, R7)
Service Provider PE routers – PE1, PE2, PE3, PE4 (Simulation R1, R2, R3, R4)
Service Provider Route Reflectors – RR8, RR9 (Simulation R8, R9)
Service Provider management router – CE1 (Simulation R10)
Customer A's routers – CE12, CE4 (Simulation R11, R15)
Customer B's routers – CE2, CE3, CE34 (Simulation R12, R13, R14)

Leave CDP enabled on all ethernet interfaces. Use this to confirm which router interfaces are actually connected together.

Service Provider Ethernet links:
Use dot1q tagged interfaces, using dot1q tag 1 for the inter P and PE links.

Customer to Provider Ethernet links:
Use dot1q tagged interfaces. Each lab will give details of the dot1 tag number to use.

Workgroup Router Roles:
Each workgroup has two points of presence (POPs), North, and South.

North POP Routers: PE1, PE2, P5, P7, RR8
North POP connected customer routers: CE2, CE12
Provider's management router (located in North POP): CE1

South POP Routers: PE3, PE4, P6, RR9
South POP connected customer routers: CE3, CE4, CE34

Service Provider Address Space Plan
All P and PE loopback addresses are taken from 172.16.0.0/24
All linknets between P and PE routers are taken from the reaminder of 172.16/12

See the lab network diagram for details of linknet addresses between routers.

NOTE: The above addressing scheme has been selected for ease of use in the labs. It does not optimize the use of the address space.

Step 2: Configuring the service provider IGP.

On all core routers (PE1, PE2, PE3, PE4, P5, P6 and P7), enable the OSPF routing process, using 1 as the autonomous system (AS) number.

Ensure that the service provider networks are configured and are being advertised by the OSPF process.

Step 3: Verification.

On each core router, verify that the OSPF router process is active.

On each core router, verify that the OSPF routing is enabled on all ethernet interfaces.

On each core router, verify that the loopback interfaces of all P and PE routers are displayed in the IP routing table.

Familiarize yourself with the routers assigned to you and the IGP setup.

Make sure you understand the network paths between the PEs and the IP addressing scheme being used.

Step 4: Solution

Ask the instructors to verify your setup once you have completed this lab.

Conclusion: Conclusion

Congratulations on completing this lab. Feel free to experiment with what you have learned, or continue to work on the devices.

Remember to save your work!

LAB 2: Establishing the Core MPLS Environment

Introduction: Introduction

Welcome to MPLS Lab 2: Establishing the Core MPLS Environment. Please progress through the steps in this lab, performing the configuration and troubleshooting exercises found in each step. A topology for the lab is supplied in order to give you a visual representation of the devices you are working with. Use the supplied interface to interact with the devices and complete the lab.

Step 1: Command list.

access-list access-list-number {permit | deny} {type-code wild-mask | address mask}
no access-list access-list-number {permit | deny} {type-code wild-mask | address mask} - - To configure the access list mechanism for filtering frames by protocol type or vendor code, use the access-list global configuration command. To remove the single specified entry from the access list, use the no form of this command.

ip cef - To enable Cisco Express Forwarding (CEF) on the route processor (RP) card, use the ip cef command in global configuration mode. To disable CEF, use the no form of this command.

mpls ip
no mpls ip - To enable MPLS forwarding of IP version 4 (IPv4) packets along normally routed paths for the platform, the mpls ip command can be used in global configuration mode (for traffic engineering [TE]) but must be used at the interface configuration mode for Label Distribution Protocol (LDP) to become active. To disable this feature, use the no form of this command.

mpls ip propagate-ttl
no mpls ip propagate-ttl [forwarded | local] - To control the generation of the time-to-live (TTL) field in the MPLS header when labels are first added to an IP packet, use the mpls ip propagate-ttl global configuration command. To use a fixed TTL value (255) for the first label of the IP packet, use the no form of this command.

mpls label protocol {ldp | tdp | both }
[no] mpls label protocol - To specify the label distribution protocol to be used on a given interface, use the mpls label protocol interface configuration command. Use the no form of the command to disable this feature.

show mpls interfaces [interface] [detail] - To display information about one or more interfaces that have been configured for label switching, use the show mpls interfaces privileged EXEC command.

show mpls ldp discovery - To display the status of the LDP discovery process, use the show mpls ldp discovery privileged EXEC command. This command generates a list of interfaces over which the LDP discovery process is running.

show mpls ldp neighbor [address | interface] [detail] - To display the status of LDP sessions, issue the show mpls ldp neighbor privileged EXEC command.

show mpls ldp bindings [network {mask | length} [longer-prefixes]] [local-label label [- label]} [remote-label label [- label] [neighbor address] [local] - To display the contents of the label information base (LIB), use the show mpls ldp bindings privileged EXEC command.

mpls ldp advertise-labels [for prefix-access-list [to peer-access-list]]
no mpls ldp advertise-labels [for prefix-access-list [to peer-access-list]] - To control the distribution of locally assigned (incoming) labels by means of LDP, use the mpls ldp advertise-labels command in global configuration mode. This command is used to control which labels are advertised to which LDP neighbors. To prevent the distribution of locally assigned labels, use the no form of this command.

Step 2: Enabling LDP on PE and P routers.

This task will establish MPLS within the service provider routing environment and will involve enabling CEF and MPLS.

On all core routers:

-Enable CEF.

-Enable LDP on the subinterface that is connected to either a P or PE router.

Step 3: Verification

You have completed this exercise when you attain these results:

-On each of the core routers, verify that the interfaces in question have been configured to use LDP.
Example

P11#sh mpls interface
Interface IP Tunnel Operational
Serial0/0.111 Yes (ldp) No Yes
Serial0/0.112 Yes (ldp) No Yes

-On each of the core routers, verify that the interface is up and has established an LDP neighbor relationship.

Example
P11#show mpls ldp discovery
Local LDP Identifier:
192.168.1.81:0
Discovery Sources:
Interfaces:
Serial0/0.111 (ldp): xmit/recv
LDP Id: 192.168.1.17:0
Serial0/0.112 (ldp): xmit/recv
LDP Id: 192.168.1.97:0

Example
P11#show mpls ldp nei
Peer LDP Ident: 192.168.1.17:0; Local LDP Ident 192.168.1.81:0
TCP connection: 192.168.1.17.646 - 192.168.1.81.11000
State: Oper; Msgs sent/rcvd: 20/23; Downstream
Up time: 00:08:03
LDP discovery sources:
Serial0/0.111, Src IP addr: 192.168.1.49
Addresses bound to peer LDP Ident:
192.168.1.17 192.168.1.49 150.1.11.18 150.1.11.34
Peer LDP Ident: 192.168.1.97:0; Local LDP Ident 192.168.1.81:0
TCP connection: 192.168.1.97.11000 - 192.168.1.81.646
State: Oper; Msgs sent/rcvd: 18/18; Downstream
Up time: 00:06:15
LDP discovery sources:
Serial0/0.112, Src IP addr: 192.168.1.114
Addresses bound to peer LDP Ident:
192.168.1.97 192.168.1.66 192.168.1.114

-On each of the core routers, verify that LDP has allocated a label for each prefix in its IP routing table.

Example
PE11#sh ip route
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, 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
192.168.1.0/24 is variably subnetted, 8 subnets, 3 masks
D 192.168.1.97/32 [90/2809856] via 192.168.1.50, 00:49:50, Serial0/0.111
D 192.168.1.112/28
[90/2681856] via 192.168.1.50, 00:49:50, Serial0/0.111
D 192.168.1.64/28 [90/3193856] via 192.168.1.50, 00:49:50, Serial0/0.111
D 192.168.1.81/32 [90/659968] via 192.168.1.50, 00:49:50, Serial0/0.111
D 192.168.1.33/32 [90/3321856] via 192.168.1.50, 00:47:00, Serial0/0.111
C 192.168.1.48/28 is directly connected, Serial0/0.111
D 192.168.1.0/24 is a summary, 00:49:20, Null0
C 192.168.1.17/32 is directly connected, Loopback0
150.1.0.0/16 is variably subnetted, 3 subnets, 2 masks
C 150.1.11.16/28 is directly connected, Serial0/0.101
D 150.1.0.0/16 is a summary, 00:49:20, Null0
C 150.1.11.32/28 is directly connected, Serial0/0.102
```

Example

```
P11#sh mpls ldp bindings
tib entry: 150.1.0.0/16, rev 16
local binding: tag: 20
remote binding: tsr: 192.168.1.17:0, tag: imp-null
remote binding: tsr: 192.168.1.97:0, tag: 20
tib entry: 150.1.11.16/28, rev 18
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 150.1.11.32/28, rev 19
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 192.168.1.0/24, rev 17
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 192.168.1.17/32, rev 14
local binding: tag: 19
remote binding: tsr: 192.168.1.17:0, tag: imp-null
remote binding: tsr: 192.168.1.97:0, tag: 19
tib entry: 192.168.1.33/32, rev 10
local binding: tag: 18
remote binding: tsr: 192.168.1.17:0, tag: 20
remote binding: tsr: 192.168.1.97:0, tag: 17
tib entry: 192.168.1.48/28, rev 12
local binding: tag: imp-null
remote binding: tsr: 192.168.1.17:0, tag: imp-null
remote binding: tsr: 192.168.1.97:0, tag: 18
tib entry: 192.168.1.64/28, rev 6
local binding: tag: 17
remote binding: tsr: 192.168.1.17:0, tag: 18
remote binding: tsr: 192.168.1.97:0, tag: imp-null
tib entry: 192.168.1.81/32, rev 8
local binding: tag: imp-null
remote binding: tsr: 192.168.1.17:0, tag: 19
remote binding: tsr: 192.168.1.97:0, tag: 16
tib entry: 192.168.1.97/32, rev 2
local binding: tag: 16
remote binding: tsr: 192.168.1.17:0, tag: 16
remote binding: tsr: 192.168.1.97:0, tag: imp-null
tib entry: 192.168.1.112/28, rev 4
local binding: tag: imp-null
remote binding: tsr: 192.168.1.17:0, tag: 17
remote binding: tsr: 192.168.1.97:0, tag: imp-null
```

- On each of the core routers, verify that LDP has received a label of the subnets and loopback interfaces of the other core routers.

Example

```
P11#sh mpls ldp bindings
tib entry: 150.1.0.0/16, rev 16
local binding: tag: 20
remote binding: tsr: 192.168.1.17:0, tag: imp-null
remote binding: tsr: 192.168.1.97:0, tag: 20
tib entry: 150.1.11.16/28, rev 18
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 150.1.11.32/28, rev 19
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 192.168.1.0/24, rev 17
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 192.168.1.17/32, rev 14
local binding: tag: 19
remote binding: tsr: 192.168.1.17:0, tag: imp-null
remote binding: tsr: 192.168.1.97:0, tag: 19
tib entry: 192.168.1.33/32, rev 10
local binding: tag: 18
remote binding: tsr: 192.168.1.17:0, tag: 20
remote binding: tsr: 192.168.1.97:0, tag: 17
tib entry: 192.168.1.48/28, rev 12
local binding: tag: imp-null
remote binding: tsr: 192.168.1.17:0, tag: imp-null
remote binding: tsr: 192.168.1.97:0, tag: 18
tib entry: 192.168.1.64/28, rev 6
local binding: tag: 17
remote binding: tsr: 192.168.1.17:0, tag: 18
remote binding: tsr: 192.168.1.97:0, tag: imp-null
tib entry: 192.168.1.81/32, rev 8
local binding: tag: imp-null
remote binding: tsr: 192.168.1.17:0, tag: 19
remote binding: tsr: 192.168.1.97:0, tag: 16
tib entry: 192.168.1.97/32, rev 2
local binding: tag: 16
remote binding: tsr: 192.168.1.17:0, tag: 16
remote binding: tsr: 192.168.1.97:0, tag: imp-null
tib entry: 192.168.1.112/28, rev 4
local binding: tag: imp-null
remote binding: tsr: 192.168.1.17:0, tag: 17
remote binding: tsr: 192.168.1.97:0, tag: imp-null
```

-Perform a traceroute from each PE router to the loopback address of the other PE router and verify that the results display the associated labels.

Example

```
Tracing the route to 192.168.1.33

1 192.168.1.50 [MPLS: Label 18 Exp 0] 164 msec 196 msec 200 msec
2 192.168.1.114 [MPLS: Label 17 Exp 0] 56 msec 56 msec 56 msec
3 192.168.1.65 40 msec 40 msec
```

Step 4: Disabling TTL propogation.

In this task, you will disable MPLS TTL propagation and verify the results.

-On each of the core routers, disable MPLS TTL propagation.

Step 5: Verification.

You have successfully completed this exercise when you attain these results:

-Perform a traceroute from a PE router to the loopback address of the other PE router and compare this display to the display obtained in the previous task.

PE11#traceroute 192.168.1.33
Type escape sequence to abort.
Tracing the route to 192.168.1.33

1 192.168.1.65 40 msec 40 msec *

NOTE: When you are troubleshooting, it may become necessary to view the core routes when doing traces. If so, it will be necessary to re-enable TTL propagation. Doing so may affect the results of the traces shown in the lab exercise verification because additional hops and labs will be displayed.

Step 6: Configuring conditional label distribution.

For the label binding displays that you did in the above task, you can see that a label is assigned to every prefix that is in the IP routing table of a router. This label assignment results in wasted label space and resources necessary to build unused label switched paths (LSPs). In this task you will use conditional label advertising to restrict the distribution of labels related to the WAN interfaces in the core.

-On each PE router, display the LSPs that are being built.

Example
PE11#sh mpls for
Local Outgoing Prefix Bytes tag Outgoing Next Hop
tag tag or VC or Tunnel Id switched interface
16 16 192.168.1.97/32 0 Se0/0.111 point2point
17 Pop tag 192.168.1.112/28 0 Se0/0.111 point2point
18 17 192.168.1.64/28 0 Se0/0.111 point2point
19 Pop tag 192.168.1.81/32 0 Se0/0.111 point2point
20 18 192.168.1.33/32 0 Se0/0.111 point2point

Note: Note that an LSP has been built to the WAN interface that connects the other PE and P router. This LSP will never be used because traffic will not normally terminate at this point.

-On your all of the core routers, configure conditional label distribution to allow only the distribution of labels related to the core loopback addresses and the interfaces that provide direct customer support.

Step 7: Verification.

You have completed this exercise when you attain these results:

-On each PE router, display the LSPs that are being built.

Example
PE11#sh mpls f
Local Outgoing Prefix Bytes tag Outgoing Next Hop
tag tag or VC or Tunnel Id switched interface
16 16 192.168.1.97/32 0 Se0/0.111 point2point
17 Untagged 192.168.1.112/28 0 Se0/0.111 point2point
18 Untagged 192.168.1.64/28 0 Se0/0.111 point2point

19 Pop tag 192.168.1.81/32 0 Se0/0.111 point2point
20 18 192.168.1.33/32 0 Se0/0.111 point2point

NOTE: An LSP is no longer built to the WAN interface that connects the other PE and P routers.

-On each P router display the LDP bindings.

Example

```
P11#sh mpls ldp bind
tib entry: 150.1.0.0/16, rev 31
local binding: tag: 20
remote binding: tsr: 192.168.1.97:0, tag: 20
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 150.1.11.16/28, rev 36
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 150.1.11.32/28, rev 37
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 192.168.1.17/32, rev 35
local binding: tag: 19
remote binding: tsr: 192.168.1.97:0, tag: 19
remote binding: tsr: 192.168.1.17:0, tag: imp-null
tib entry: 192.168.1.33/32, rev 32
local binding: tag: 18
remote binding: tsr: 192.168.1.97:0, tag: 17
remote binding: tsr: 192.168.1.17:0, tag: 20
tib entry: 192.168.1.48/28, rev 26
local binding: tag: imp-null
tib entry: 192.168.1.64/28, rev 27
local binding: tag: 17
tib entry: 192.168.1.81/32, rev 34
local binding: tag: imp-null
remote binding: tsr: 192.168.1.97:0, tag: 16
remote binding: tsr: 192.168.1.17:0, tag: 19
tib entry: 192.168.1.97/32, rev 33
local binding: tag: 16
remote binding: tsr: 192.168.1.97:0, tag: imp-null
remote binding: tsr: 192.168.1.17:0, tag: 16
tib entry: 192.168.1.112/28, rev 30
local binding: tag: imp-null
```

NOTE: The prefix assigned to the WAN interface connecting the other P and PE routers no longer has a remote label assigned. Further, none of the core WAN interfaces have remote labels assigned. This lessening of assignments results in a reduced label space, which saves memory resources.

Step 8: Solution.

The solution to this lab can be found at

Conclusion: Conclusion

Congratulations on completing this lab. Feel free to experiment with what you have learned, or continue to work on the devices.