

CCIE Service Provider Lab Workbook v4.0 (<http://labs.ine.com/workbook/toc/service-provider-v4>) » CCIE SP v4 Advanced Technology Labs - MPLS

› LDP IGP Synchronization

CONTENTS

« LDP Label Allocation Filtering (</workbook/view/service-provider-v4/task/ldp-label-allocation-filtering-Mjg1MA%3D%3D>) | LDP Session Protection (</workbook/view/service-provider-v4/task/ldp-session-protection-Mjg1Mg%3D%3D>) »

Last updated: April 23, 2016

Note:

Initial Configuration & Diagrams: Load the initial configuration files for the section named **Basic L3VPN**, which can be found in [CCIE SPv4 Topology Diagrams & Initial Configurations](http://labs.ine.com/workbook/view/service-provider-v4/task/ccie-spv4-topology-diagrams-initial-configs) (<http://labs.ine.com/workbook/view/service-provider-v4/task/ccie-spv4-topology-diagrams-initial-configs>). Refer to the **Base IPv4 Diagram** in order to complete this task.

Task

- Configure MPLS Label Distribution with LDP on all links connecting R2, R3, R4, R5, R6, and XR1.
- R2 and XR1 are preconfigured as PE routers for the MPLS L3VPN customer routers R1 and XR2; at this point R1 and XR2 should have reachability to each other's Loopback0 networks.
- In the core of the SP network change the OSPF cost of every transit link to 100 with the exception of the links between R2 & R3, R3 & R6, and R6 & XR1.
- Configure LDP and IGP Synchronization with OSPF on all routers in the core of the SP network.
- To test this, filter out all LDP traffic that R6 is receiving and ensure that labeled traffic through the core reroutes around the unsynchronized links of R6.

Configuration Click to collapse

```
R2:
interface GigabitEthernet1.24
 ip ospf cost 100
!
router ospf 1
 mpls ldp autoconfig area 0
 mpls ldp sync
```

```
R3:
interface GigabitEthernet1.34
 ip ospf cost 100
!
router ospf 1
 mpls ldp autoconfig area 0
 mpls ldp sync
```

```
R4:
interface GigabitEthernet1.24
 ip ospf cost 100
!
interface GigabitEthernet1.34
 ip ospf cost 100
!
interface GigabitEthernet1.45
 ip ospf cost 100
!
interface GigabitEthernet1.46
 ip ospf cost 100
!
router ospf 1
 mpls ldp autoconfig area 0
 mpls ldp sync
```

```
R5:
interface GigabitEthernet1.45
 ip ospf cost 100
!
interface GigabitEthernet1.56
 ip ospf cost 100
!
interface GigabitEthernet1.519
 ip ospf cost 100
!
router ospf 1
 mpls ldp autoconfig area 0
 mpls ldp sync
```

```
R6:
interface GigabitEthernet1.46
 ip ospf cost 100
!
interface GigabitEthernet1.56
 ip ospf cost 100
```

```

!
router ospf 1
  mpls ldp autoconfig area 0
  mpls ldp sync

XR1:
mpls ldp
  router-id 19.19.19.19
!
router ospf 1
  mpls ldp sync
  mpls ldp auto-config
  area 0
  interface GigabitEthernet0/0/0.519
  cost 100

```

Verification

MPLS LDP IGP Synchronization is used to prevent traffic blackholes in the core of the MPLS network when an error in LDP configuration or operation causes IGP to attempt to route labeled traffic over a non-labeled path. LDP Sync prevents these blackholes by configuring the IGP process (either OSPF or IS-IS) to advertise the highest possible cost values for links that do not have their LDP adjacencies established and properly converged.

The idea is that if LDP fails on an interface, the router will begin advertising a very high IGP cost for that link, which ideally should cause the IGP network to reroute and avoid the link on which LDP is broken. This feature also helps prevent blackholes when a router reloads. If IGP comes up before LDP has fully exchanged labels, then traffic could end up being blackholed.

To see this feature in action in this particular scenario, we first need to see which way the end customer traffic is routing through the core of the Service Provider network. This can be seen through the following traceroute output.

```

R1#traceroute 20.20.20.20 source 1.1.1.1
Type escape sequence to abort.
Tracing the route to 20.20.20.20
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.2.2 4 msec 1 msec 1 msec
 2 20.2.3.3 [MPLS: Labels 20/16007 Exp 0] 12 msec 8 msec 9 msec
 3 20.3.6.6 [MPLS: Labels 20/16007 Exp 0] 18 msec 31 msec 31 msec
 4 20.6.19.19 22 msec 14 msec 15 msec
 5 10.19.20.20 15 msec * 10 msec

RP/0/3/CPU0:XR2#traceroute 1.1.1.1 source 20.20.20.20
Fri May 1 23:52:14.596 UTC
Type escape sequence to abort.
Tracing the route to 1.1.1.1
 1 10.19.20.19 9 msec 0 msec 0 msec
 2 20.6.19.6 [MPLS: Labels 16/31 Exp 0] 29 msec 9 msec 0 msec
 3 20.3.6.3 [MPLS: Labels 16/31 Exp 0] 0 msec 9 msec 9 msec
 4 10.1.2.2 [MPLS: Label 31 Exp 0] 0 msec 0 msec 0 msec
 5 10.1.2.1 9 msec * 9 msec

```

For both traffic from R1 to XR2 and back, the links between R3 and R6 are preferred in the transit path of the core of the network. This is due to the fact that all other links in the core of the OSPF network have their cost raised to 100, as seen below.

CONTENTS

R2#show ip ospf interface brief

| Interface | PID | Area | IP Address/Mask | Cost | State | Nbrs | F/C |
|-----------|-----|------|-----------------|------|-------|------|-----|
| Lo0 | 1 | 0 | 2.2.2.2/32 | 1 | LOOP | 0/0 | |
| Gi1.24 | 1 | 0 | 20.2.4.2/24 | 100 | BDR | 1/1 | |
| Gi1.23 | 1 | 0 | 20.2.3.2/24 | 1 | BDR | 1/1 | |

R3#show ip ospf interface brief

| Interface | PID | Area | IP Address/Mask | Cost | State | Nbrs | F/C |
|-----------|-----|------|-----------------|------|-------|------|-----|
| Lo0 | 1 | 0 | 3.3.3.3/32 | 1 | LOOP | 0/0 | |
| Gi1.36 | 1 | 0 | 20.3.6.3/24 | 1 | BDR | 1/1 | |
| Gi1.34 | 1 | 0 | 20.3.4.3/24 | 100 | BDR | 1/1 | |
| Gi1.23 | 1 | 0 | 20.2.3.3/24 | 1 | DR | 1/1 | |

R4#show ip ospf interface brief

| Interface | PID | Area | IP Address/Mask | Cost | State | Nbrs | F/C |
|-----------|-----|------|-----------------|------|-------|------|-----|
| Lo0 | 1 | 0 | 4.4.4.4/32 | 1 | LOOP | 0/0 | |
| Gi1.46 | 1 | 0 | 20.4.6.4/24 | 100 | BDR | 1/1 | |
| Gi1.45 | 1 | 0 | 20.4.5.4/24 | 100 | BDR | 1/1 | |
| Gi1.34 | 1 | 0 | 20.3.4.4/24 | 100 | DR | 1/1 | |
| Gi1.24 | 1 | 0 | 20.2.4.4/24 | 100 | DR | 1/1 | |

R5#show ip ospf interface brief

| Interface | PID | Area | IP Address/Mask | Cost | State | Nbrs | F/C |
|-----------|-----|------|-----------------|------|-------|------|-----|
| Lo0 | 1 | 0 | 5.5.5.5/32 | 1 | LOOP | 0/0 | |
| Gi1.519 | 1 | 0 | 20.5.19.5/24 | 100 | BDR | 1/1 | |
| Gi1.56 | 1 | 0 | 20.5.6.5/24 | 100 | BDR | 1/1 | |
| Gi1.45 | 1 | 0 | 20.4.5.5/24 | 100 | DR | 1/1 | |

R6#show ip ospf interface brief

| Interface | PID | Area | IP Address/Mask | Cost | State | Nbrs | F/C |
|-----------|-----|------|-----------------|------|-------|------|-----|
| Lo0 | 1 | 0 | 6.6.6.6/32 | 1 | LOOP | 0/0 | |
| Gi1.619 | 1 | 0 | 20.6.19.6/24 | 1 | BDR | 1/1 | |
| Gi1.56 | 1 | 0 | 20.5.6.6/24 | 100 | DR | 1/1 | |
| Gi1.46 | 1 | 0 | 20.4.6.6/24 | 100 | DR | 1/1 | |
| Gi1.36 | 1 | 0 | 20.3.6.6/24 | 1 | DR | 1/1 | |

RP/0/0/CPU0:XR1#show ip ospf interface brief

Fri May 1 23:56:25.278 UTC

* Indicates MADJ interface, (P) Indicates fast detect hold down state

Interfaces for OSPF 1

| Interface | PID | Area | IP Address/Mask | Cost | State | Nbrs | F/C |
|-------------|-----|------|-----------------|------|-------|------|-----|
| Lo0 | 1 | 0 | 19.19.19.19/32 | 1 | LOOP | 0/0 | |
| Gi0/0/0.519 | 1 | 0 | 20.5.19.19/24 | 100 | DR | 1/1 | |
| Gi0/0/0.619 | 1 | 0 | 20.6.19.19/24 | 1 | DR | 1/1 | |

Additionally at this point LDP IGP Sync is enabled, and all interfaces running OSPF also have LDP enabled, since LDP autoconfig was enabled. This means that the IGP and LDP domains should be synchronized, as seen below.

CONTENTS >

```
R2#show mpls ldp igp sync
```

```
GigabitEthernet1.23:
```

```
LDP configured; LDP-IGP Synchronization enabled.  
Sync status: sync achieved; peer reachable.  
Sync delay time: 0 seconds (0 seconds left)  
IGP holddown time: infinite.  
Peer LDP Ident: 3.3.3.3:0  
IGP enabled: OSPF 1
```

```
GigabitEthernet1.24:
```

```
LDP configured; LDP-IGP Synchronization enabled.  
Sync status: sync achieved; peer reachable.  
Sync delay time: 0 seconds (0 seconds left)  
IGP holddown time: infinite.  
Peer LDP Ident: 4.4.4.4:0  
IGP enabled: OSPF 1
```

```
R3#show mpls ldp igp sync
```

```
GigabitEthernet1.23:
```

```
LDP configured; LDP-IGP Synchronization enabled.  
Sync status: sync achieved; peer reachable.  
Sync delay time: 0 seconds (0 seconds left)  
IGP holddown time: infinite.  
Peer LDP Ident: 2.2.2.2:0  
IGP enabled: OSPF 1
```

```
GigabitEthernet1.34:
```

```
LDP configured; LDP-IGP Synchronization enabled.  
Sync status: sync achieved; peer reachable.  
Sync delay time: 0 seconds (0 seconds left)  
IGP holddown time: infinite.  
Peer LDP Ident: 4.4.4.4:0  
IGP enabled: OSPF 1
```

```
GigabitEthernet1.36:
```

```
LDP configured; LDP-IGP Synchronization enabled.  
Sync status: sync achieved; peer reachable.  
Sync delay time: 0 seconds (0 seconds left)  
IGP holddown time: infinite.  
Peer LDP Ident: 6.6.6.6:0  
IGP enabled: OSPF 1
```

```
R4#show mpls ldp igp sync
```

```
GigabitEthernet1.24:
```

```
LDP configured; LDP-IGP Synchronization enabled.  
Sync status: sync achieved; peer reachable.  
Sync delay time: 0 seconds (0 seconds left)  
IGP holddown time: infinite.  
Peer LDP Ident: 2.2.2.2:0  
IGP enabled: OSPF 1
```

```
GigabitEthernet1.34:
```

```
LDP configured; LDP-IGP Synchronization enabled.  
Sync status: sync achieved; peer reachable.  
Sync delay time: 0 seconds (0 seconds left)  
IGP holddown time: infinite.  
Peer LDP Ident: 3.3.3.3:0
```

```
IGP enabled: OSPF 1
GigabitEthernet1.45:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite.
  Peer LDP Ident: 5.5.5.5:0
  IGP enabled: OSPF 1
GigabitEthernet1.46:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite.
  Peer LDP Ident: 6.6.6.6:0
  IGP enabled: OSPF 1
```

```
R5#show mpls ldp igp sync
```

```
GigabitEthernet1.45:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite.
  Peer LDP Ident: 4.4.4.4:0
  IGP enabled: OSPF 1
GigabitEthernet1.56:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite.
  Peer LDP Ident: 6.6.6.6:0
  IGP enabled: OSPF 1
GigabitEthernet1.519:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite.
  Peer LDP Ident: 19.19.19.19:0
  IGP enabled: OSPF 1
```

```
R6#show mpls ldp igp sync
```

```
GigabitEthernet1.36:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite.
  Peer LDP Ident: 3.3.3.3:0
  IGP enabled: OSPF 1
GigabitEthernet1.46:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite
  Peer LDP Ident: 4.4.4.4:0
  IGP enabled: OSPF 1
```

```
GigabitEthernet1.56:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite.
  Peer LDP Ident: 5.5.5.5:0
  IGP enabled: OSPF 1
```

```
GigabitEthernet1.619:
  LDP configured; LDP-IGP Synchronization enabled.
  Sync status: sync achieved; peer reachable.
  Sync delay time: 0 seconds (0 seconds left)
  IGP holddown time: infinite.
  Peer LDP Ident: 19.19.19.19:0
  IGP enabled: OSPF 1
```

```
RP/0/0/CPU0:XR1#show mpls ldp igp sync
```

```
Fri May 1 23:58:02.882 UTC
```

```
GigabitEthernet0/0/0.519:
  VRF: 'default' (0x60000000)
  Sync delay: Disabled
  Sync status: Ready
  Peers:
    5.5.5.5:0
```

```
GigabitEthernet0/0/0.619:
  VRF: 'default' (0x60000000)
  Sync delay: Disabled
  Sync status: Ready
  Peers:
    6.6.6.6:0
```

Since synchronization has been achieved it means that OSPF will be advertising the normal cost value for its connected links. Below we see that R6 is advertising a cost of 1 for the links to R3 and XR1, while advertising a cost of 100 on the links to R4 and R5.

```
RP/0/0/CPU0:XR1#show ospf database router 6.6.6.6
```

```
Sat May 2 00:00:36.841 UTC
```

```
OSPF Router with ID (19.19.19.19) (Process ID 1)
```

```
Router Link States (Area 0)
```

```
Routing Bit Set on this LSA
```

```
LS age: 551
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: Router Links
```

```
Link State ID: 6.6.6.6
```

```
Advertising Router: 6.6.6.6
```

```
LS Seq Number: 80000007
```

```
Checksum: 0x96f9
```

```
Length: 84
```

```
Number of Links: 5
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 6.6.6.6
```

```
(Link Data) Network Mask: 255.255.255.255
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 1
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 20.6.19.6
```

```
(Link Data) Router Interface address: 20.6.19.6
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 1
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 20.5.6.6
```

```
(Link Data) Router Interface address: 20.5.6.6
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 100
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 20.4.6.6
```

```
(Link Data) Router Interface address: 20.4.6.6
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 100
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 20.3.6.6
```

```
(Link Data) Router Interface address: 20.3.6.6
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 1
```

LDP IGP Sync is now protecting the network against a failure of LDP that would normally blackhole traffic. For example suppose that R6 has an access-list configured to filter traffic in the data plane that arrives inbound on its local interfaces. Additionally, someone misconfigures this filter so that LDP traffic (TCP port 646) is dropped. Since the LDP holdtime is very long (180 seconds by default), they wouldn't notice that the problems occurs immediately. Instead about 3 minutes after the filter is configured the

LDP adjacencies of R6 would drop, causing traffic to blackhole. However with LDP IGP Sync on, the IGP process of R6 would detect that LDP Synchronization has been lost, and would start advertising a very high cost for its local links in the attempt to reroute traffic around its links and/or node.

Below we see an ACL configured on R6 to filter out the LDP adjacencies. This filter is applied in on all interfaces at around time index 00:04:40.236.

```
R6#config t
Enter configuration commands, one per line. End with CNTL/Z.
R6(config)#access-list 100 deny tcp any any eq 646
R6(config)#access-list 100 deny tcp any eq 646 any
R6(config)#access-list 100 permit ip any any
R6(config)#int Gig1.36
R6(config-subif)#ip access-group 100 in
R6(config-subif)#int Gig1.46
R6(config-subif)#ip access-group 100 in
R6(config-subif)#int Gig1.56
R6(config-subif)#ip access-group 100 in
R6(config-subif)#int Gig1.619
R6(config-subif)#ip access-group 100 in
R6(config)#end
R6#
*May 2 00:04:40.236: %SYS-5-CONFIG_I: Configured from console by console

R6#show mpls ldp neighbor 3.3.3.3 detail

Peer LDP Ident: 3.3.3.3:0; Local LDP Ident 6.6.6.6:0

TCP connection: 3.3.3.3.646 - 6.6.6.6.46449

Password: not required, none, in use

State: Oper; Msgs sent/rcvd: 52/52; Downstream; Last TIB rev sent 30

Up time: 00:29:50; UID: 27; Peer Id 0

LDP discovery sources:

GigabitEthernet1.36; Src IP addr: 20.3.6.3

holdtime: 15000 ms, hello interval: 5000 ms

Addresses bound to peer LDP Ident:

3.3.3.3      20.2.3.3    20.3.4.3    20.3.6.3

Peer holdtime: 180000 ms; KA interval: 60000 ms; Peer state: estab

NSR: Not Ready

Capabilities Sent:

[ICCP (type 0x0405) MajVer 1 MinVer 0]
[Dynamic Announcement (0x0506)]
[mLDP Point-to-Multipoint (0x0508)]
[mLDP Multipoint-to-Multipoint (0x0509)]
[Typed Wildcard (0x050B)]

Capabilities Received:

[ICCP (type 0x0405) MajVer 1 MinVer 0]
[Dynamic Announcement (0x0506)]
[mLDP Point-to-Multipoint (0x0508)]
[mLDP Multipoint-to-Multipoint (0x0509)]
[Typed Wildcard (0x050B)]
```

At this point the LDP adjacencies of R6 are still up, so traffic is routed through the normal lowest cost path in the core.

```
R1#traceroute 20.20.20.20 source 1.1.1.1
```

```
Type escape sequence to abort.
```

```
Tracing the route to 20.20.20.20
```

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

```
 1 10.1.2.2 4 msec 1 msec 1 msec
 2 20.2.3.3 [MPLS: Labels 20/16007 Exp 0] 9 msec 6 msec 6 msec
 3 20.3.6.6 [MPLS: Labels 20/16007 Exp 0] 27 msec 31 msec 37 msec
 4 20.6.19.19 22 msec 15 msec 15 msec
 5 10.19.20.20 16 msec * 12 msec
```

```
R2#show ip route 19.19.19.19
```

```
Routing entry for 19.19.19.19/32
```

```
Known via "ospf 1", distance 110, metric 4, type intra area
```

```
Last update from 20.2.3.3 on GigabitEthernet1.23, 00:14:55 ago
```

```
Routing Descriptor Blocks:
```

```
* 20.2.3.3, from 19.19.19.19, 00:14:55 ago, via GigabitEthernet1.23
```

```
Route metric is 4, traffic share count is 1
```

About 3 minutes later R6 starts to lose its LDP adjacencies.

```
R6#
```

```
*May 2 00:06:19.645: %LDP-5-NBRCHG: LDP Neighbor 19.19.19.19:0 (4) is DOWN (Session KeepAlive Timer expired)
```

```
R6#
```

```
*May 2 00:06:37.089: %LDP-5-NBRCHG: LDP Neighbor 5.5.5.5:0 (2) is DOWN (Session KeepAlive Timer expired)
```

```
R6#
```

```
*May 2 00:07:25.856: %LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (3) is DOWN (Session KeepAlive Timer expired)
```

```
R6#
```

```
*May 2 00:07:32.016: %LDP-5-NBRCHG: LDP Neighbor 3.3.3.3:0 (1) is DOWN (Session KeepAlive Timer expired)
```

This causes IGP Sync to be lost.

```
R6#show mpls ldp igp sync
```

```
GigabitEthernet1.36:
```

```
LDP configured; LDP-IGP Synchronization enabled.
```

```
Sync status: sync not achieved; peer reachable.
```

```
Sync delay time: 0 seconds (0 seconds left)
```

```
IGP holddown time: infinite.
```

```
IGP enabled: OSPF 1
```

```
GigabitEthernet1.46:
```

```
LDP configured; LDP-IGP Synchronization enabled.
```

```
Sync status: sync not achieved; peer reachable.
```

```
Sync delay time: 0 seconds (0 seconds left)
```

```
IGP holddown time: infinite.
```

```
IGP enabled: OSPF 1
```

```
GigabitEthernet1.56:
```

```
LDP configured; LDP-IGP Synchronization enabled.
```

```
Sync status: sync not achieved; peer reachable.
```

```
Sync delay time: 0 seconds (0 seconds left)
```

```
IGP holddown time: infinite.
```

```
IGP enabled: OSPF 1
```

```
GigabitEthernet1.619:
```

```
LDP configured; LDP-IGP Synchronization enabled.
```

```
Sync status: sync not achieved; peer reachable.
```

```
Sync delay time: 0 seconds (0 seconds left)
```

```
IGP holddown time: infinite.
```

```
IGP enabled: OSPF 1
```

R6 now advertises a high cost for its connected links.

```
RP/0/0/CPU0:XR1#show ospf database router 6.6.6.6
```

```
Mon Mar 26 15:05:54.541 UTC
```

```
OSPF Router with ID (19.19.19.19) (Process ID 1)
```

```
Router Link States (Area 0)
```

```
Routing Bit Set on this LSA LS age: 366
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: Router Links
```

```
Link State ID: 6.6.6.6
```

```
Advertising Router: 6.6.6.6
```

```
LS Seq Number: 8000000b
```

```
Checksum: 0x88ce
```

```
Length: 84
```

```
Number of Links: 5
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 6.6.6.6
```

```
(Link Data) Network Mask: 255.255.255.255
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 1
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 20.6.19.6
```

```
(Link Data) Router Interface address: 20.6.19.6
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 65535
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 20.5.6.6
```

```
(Link Data) Router Interface address: 20.5.6.6
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 65535
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 20.4.6.6
```

```
(Link Data) Router Interface address: 20.4.6.6
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 65535
```

```
Link connected to: a Transit Network
```

```
(Link ID) Designated Router address: 20.3.6.6
```

```
(Link Data) Router Interface address: 20.3.6.6
```

```
Number of TOS metrics: 0
```

```
TOS 0 Metrics: 65535
```

The final result is that the end customer's traffic is transparently rerouted around the LDP failure.

```
R1#traceroute 20.20.20.20 source 1.1.1.1
```

```
Type escape sequence to abort.
```

```
Tracing the route to 20.20.20.20
```

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

```
 1 10.1.1.2.2 1 msec 2 msec 0 msec
 2 20.2.4.4 [MPLS: Labels 20/16007 Exp 0] 10 msec 8 msec 8 msec
 3 20.4.5.5 [MPLS: Labels 20/16007 Exp 0] 20 msec 31 msec 31 msec
 4 20.5.19.19 21 msec 16 msec 13 msec
 5 10.19.20.20 15 msec * 10 msec
```

```
R2#show ip route 19.19.19.19
```

```
Routing entry for 19.19.19.19/32
```

```
Known via "ospf 1", distance 110, metric 301, type intra area
```

```
Last update from 20.2.4.4 on GigabitEthernet1.24, 00:02:37 ago
```

```
Routing Descriptor Blocks:
```

```
* 20.2.4.4, from 19.19.19.19, 00:02:37 ago, via GigabitEthernet1.24
```

```
Route metric is 301, traffic share count is 1
```

Now let's examine the problem *without* LDP IGP Sync. All routers remove the LDP IGP Sync feature, and R6 removes its filter that is breaking LDP.

```
R2#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)#router ospf 1
R2(config-router)#no mpls ldp sync
R2(config-router)#end
R2#

R3#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R3(config)#router ospf 1
R3(config-router)#no mpls ldp sync
R3(config-router)#end
R3#

R4#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R4(config)#router ospf 1
R4(config-router)#no mpls ldp sync
R4(config-router)#end
R4#

R5#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R5(config)#router ospf 1
R5(config-router)#no mpls ldp sync
R5(config-router)#end
R5#

R6#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R6(config)#no access-list 100
R6(config)#router ospf 1
R6(config-router)#no mpls ldp sync
R6(config-router)#end
00:41:22: %SYS-5-CONFIG_I: Configured from console by console
00:41:23: %LDP-5-NBRCHG: LDP Neighbor 3.3.3.3:0 (1) is UP
00:41:23: %LDP-5-NBRCHG: LDP Neighbor 5.5.5.5:0 (2) is UP
00:41:24: %LDP-5-NBRCHG: LDP Neighbor 19.19.19.19:0 (3) is UP
00:41:27: %LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (4) is UP

RP/0/0/CPU0:XR1#config t
Mon Mar 26 15:10:30.219 UTC
RP/0/0/CPU0:XR1(config)#router ospf 1
RP/0/0/CPU0:XR1(config-ospf)#no mpls ldp sync
RP/0/0/CPU0:XR1(config-ospf)#commit
Sat May 2 00:11:45.865 UTC
```

At this point all LDP adjacencies are working and traffic is routing as normal via the R2 > R3 > R6 > XR1 path and back.

```
R1#traceroute 20.20.20.20 source 1.1.1.1
Type escape sequence to abort.
Tracing the route to 20.20.20.20
VRF info: (vrf in name/id, vrf out name/id)
  1 10.1.1.2.2 4 msec 1 msec 1 msec
  2 20.2.3.3 [MPLS: Labels 20/16007 Exp 0] 6 msec 4 msec 7 msec
  3 20.3.6.6 [MPLS: Labels 20/16007 Exp 0] 30 msec 31 msec 31 msec
  4 20.6.19.19 20 msec 16 msec 15 msec
  5 10.19.20.20 14 msec * 8 msec
```

Now R6 configures its ACL filter again which breaks LDP. The ACL is still applied from before.

```
R6#config t
Enter configuration commands, one per line. End with CNTL/Z.
R6(config)#access-list 100 deny tcp any any eq 646
R6(config)#access-list 100 deny tcp any eq 646 any
*May 2 00:14:07.832: %LDP-5-NBRCHG: LDP Neighbor 19.19.19.19:0 (4) is DOWN (Discovery Hello Hold Timer expired)
*May 2 00:14:08.137: %LDP-5-NBRCHG: LDP Neighbor 3.3.3.3:0 (1) is DOWN (Discovery Hello Hold Timer expired)
R6(config)#access-list 100 permit ip any any
R6(config)#
*May 2 00:14:10.304: %LDP-5-NBRCHG: LDP Neighbor 5.5.5.5:0 (3) is DOWN (Discovery Hello Hold Timer expired)
*May 2 00:14:10.975: %LDP-5-NBRCHG: LDP Neighbor 4.4.4.4:0 (2) is DOWN (Discovery Hello Hold Timer expired)
```

Without LDP IGP Sync enabled, OSPF does not react to this change in the LDP topology.

```
R2#show ip route 19.19.19.19
Routing entry for 19.19.19.19/32
  Known via "ospf 1", distance 110, metric 4, type intra area
  Last update from 20.2.3.3 on GigabitEthernet1.23, 00:02:58 ago
  Routing Descriptor Blocks:
  * 20.2.3.3, from 19.19.19.19, 00:02:58 ago, via GigabitEthernet1.23
    Route metric is 4, traffic share count is 1
```

This means that traffic within the core is fine, but end customer traffic transiting the core is blackholed because of the failed Label Switch Path.

```
R2#ping 19.19.19.19 source 2.2.2.2

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 19.19.19.19, 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 = 1/3/8 ms
```

```
R1#traceroute 20.20.20.20 source 1.1.1.1
```

```
Type escape sequence to abort.

Tracing the route to 20.20.20.20
```

```
 1 10.1.1.2 4 msec 0 msec 0 msec
 2 * * *
 3 * * *
 4 * * *

<snip>
```

```
RP/0/0/CPU0:XR2#ping 1.1.1.1 source 20.20.20.20
```

```
Sat May 2 00:18:40.357 UTC
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 1.1.1.1, timeout is 2 seconds:
```

```
U.U.U
```

```
Success rate is 0 percent (0/5)
```

Although pings work between the PE's loopbacks, the LSP between them is broken. This can be further verified by using the MPLS ping utility.

```
R2#ping mpls ipv4 19.19.19.19/32 verbose source 2.2.2.2
```

```
Sending 5, 72-byte MPLS Echos to Target FEC Stack TLV descriptor,

timeout is 2 seconds, send interval is 0 msec:
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'l' - Label switched with FEC change, 'd' - see DDMAP for return code,
'X' - unknown return code, 'x' - return code 0
```

```
Type escape sequence to abort.
```

```
B size 72, reply addr 20.2.3.3, return code 9
B size 72, reply addr 20.2.3.3, return code 9
B size 72, reply addr 20.2.3.3, return code 9
B size 72, reply addr 20.2.3.3, return code 9
B size 72, reply addr 20.2.3.3, return code 9
```

```
Success rate is 0 percent (0/5)
```

```
Total Time Elapsed 30 ms
```

```
R2#traceroute mpls ipv4 19.19.19.19/32 verbose source 2.2.2.2 ttl 4
more work needed here to demux the tfs subtlv and to display the right output
```

```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'l' - Label switched with FEC change, 'd' - see DDMAP for return code,
'X' - unknown return code, 'x' - return code 0
```

Type escape sequence to abort.

```
0 20.2.3.2 20.2.3.3 MRU 1500 [Labels: 20 Exp: 0]
B 1 20.2.3.3 20.3.6.6 MRU 1500 [No Label] 14 ms, ret code 9
B 2 20.2.3.3 20.3.6.6 MRU 1500 [No Label] 7 ms, ret code 9
B 3 20.2.3.3 20.3.6.6 MRU 1500 [No Label] 2 ms, ret code 9
B 4 20.2.3.3 20.3.6.6 MRU 1500 [No Label] 2 ms, ret code 9
```

```
R3#show mpls forwarding-table 19.19.19.19 detail
```

| Local | Outgoing | Prefix | Bytes | Label | Outgoing | Next Hop |
|---|----------|----------------|----------|-------|-----------|----------|
| Label | Label | or Tunnel Id | Switched | | interface | |
| 20 | No Label | 19.19.19.19/32 | 1265 | | Gi1.36 | 20.3.6.6 |
| MAC/Encaps=18/18, MRU=1504, Label Stack{} | | | | | | |
| 0050569E5CEC0050569E6E6A810000240800 | | | | | | |
| No output feature configured | | | | | | |

« LDP Label Allocation Filtering (/workbook/view/service-provider-v4/task/ldp-label-allocation-filtering-Mjg1MA%3D%3D) | LDP Session Protection (/workbook/view/service-provider-v4/task/ldp-session-protection-Mjg1Mg%3D%3D) »