

Label Distribution Protocol and Basic MPLS Configuration

APNIC

APNIC

Issue Date: [201609]

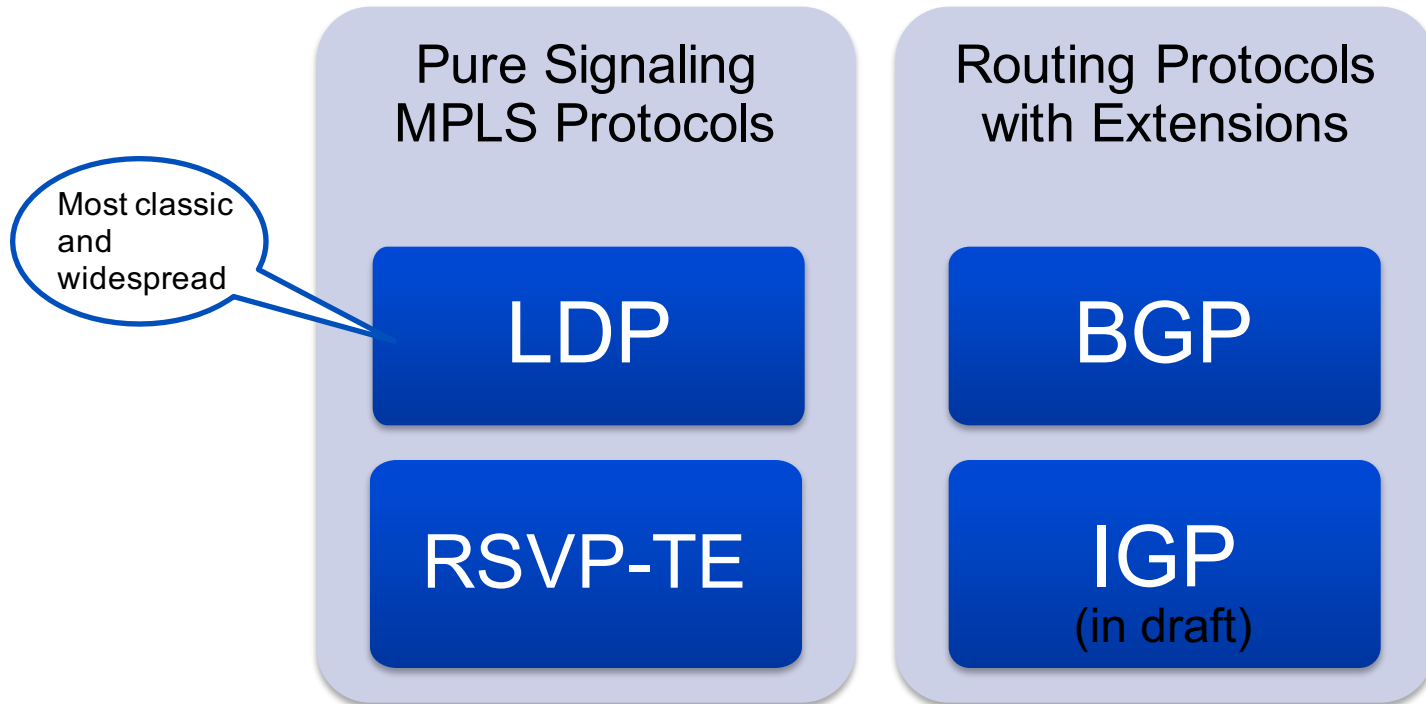
Revision: [01]



Label Distribution Protocol

MPLS Builders

Which protocols can set up Label Switched Path?



Advantages of LDP

- **Reliability**
 - LDP uses reliable TCP as the transport protocol for all but the discovery messages.
- **Auto provision**
 - Abilities to set up LSPs dynamically based on routing information
- **Plug-and-play**
 - Simple deployment and configuration
- **Support for a large number of LSPs**

LDP Identifier

- An LDP Identifier is a six octet quantity used to identify an LSR label space.

| 4 byte | 2 byte |
|------------|----------------|
| LSR ID | Label Space ID |
| 10.10.1.1 | 0 |
| 20.20.20.2 | 6 |

Label Space ID = 0
Label space is **per platform**
Label Space ID ≠ 0
Label space is **per interface**

```
R2#show mpls ldp discovery
```

```
Local LDP Identifier:
```

```
2.2.2.2:0
```

```
Discovery Sources:
```

```
Interfaces:
```

```
FastEthernet0/0 (ldp): xmit/recv
```

```
LDP Id: 3.3.3.3:0
```

```
Ethernet1/0 (ldp): xmit/recv
```

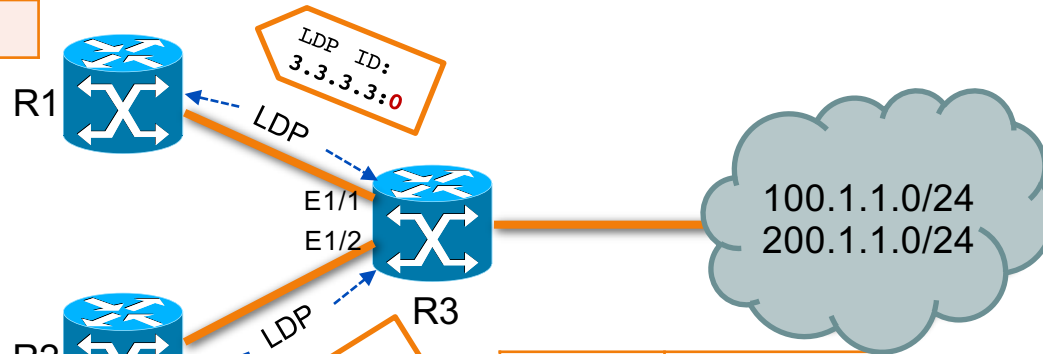
```
LDP Id: 1.1.1.1:0
```

Cisco IOS

Label Space – Per Platform

- In per-platform label space, **one single label** is assigned to a destination network and announced to all neighbors. The label must be locally **unique and valid on all incoming interfaces**.

| Prefix | Out Label |
|--------------|------------|
| 100.1.1.0/24 | 100 |
| 200.1.1.0/24 | 200 |



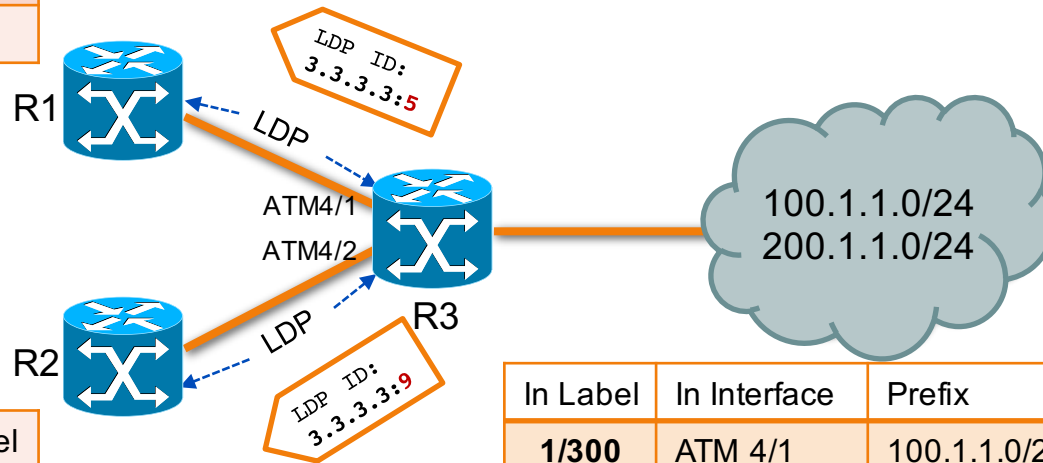
| Prefix | Out Label |
|--------------|------------|
| 100.1.1.0/24 | 100 |
| 200.1.1.0/24 | 200 |

| In Label | Prefix |
|------------|--------------|
| 100 | 100.1.1.0/24 |
| 200 | 200.1.1.0/24 |

Label Space – Per Interface

- In per-interface label space, local labels are assigned to IP destination prefixes **on a per-interface basis**. These labels must be unique on a per-interface basis.

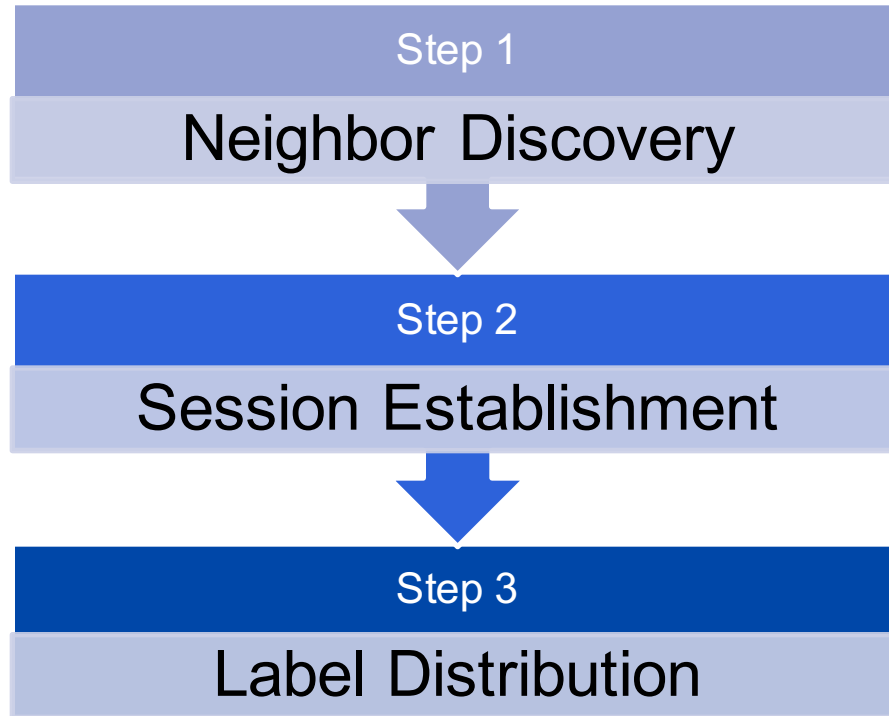
| Prefix | Out Label |
|--------------|--------------|
| 100.1.1.0/24 | 1/300 |
| 200.1.1.0/24 | 1/200 |



| Prefix | Out Label |
|--------------|--------------|
| 100.1.1.0/24 | 1/400 |
| 200.1.1.0/24 | 1/500 |

| In Label | In Interface | Prefix |
|--------------|--------------|--------------|
| 1/300 | ATM 4/1 | 100.1.1.0/24 |
| 1/200 | ATM 4/1 | 200.1.1.0/24 |
| 1/400 | ATM 4/2 | 100.1.1.0/24 |
| 1/500 | ATM 4/2 | 200.1.1.0/24 |

LDP Operations



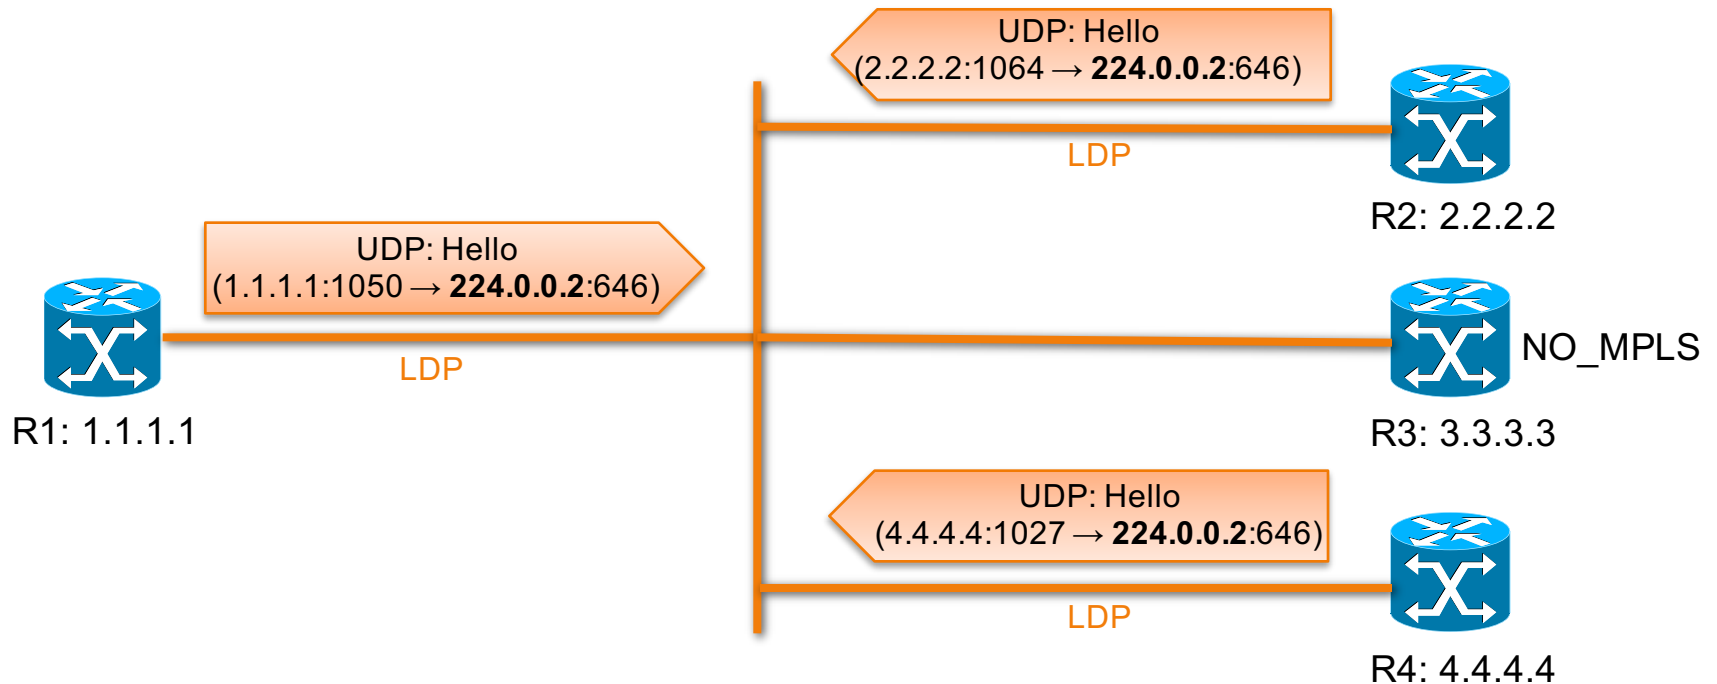
LDP Messages

| Category | Function | Message Name |
|--------------------|---|---------------------|
| Discovery | Announce and maintain the presence of an LSR in a network | Hello |
| Session | Establish, maintain, and terminate sessions between LDP peers | Initialization |
| | | Keepalive |
| Label Distribution | Create, change, and delete label mappings for FECs | Label Release |
| | | Label Request |
| | | Label Abort Request |
| | | Label Mapping |
| | | Label Withdrawal |
| Notification | Provide advisory information and to signal error information | Notification |

(Not list all the messages)

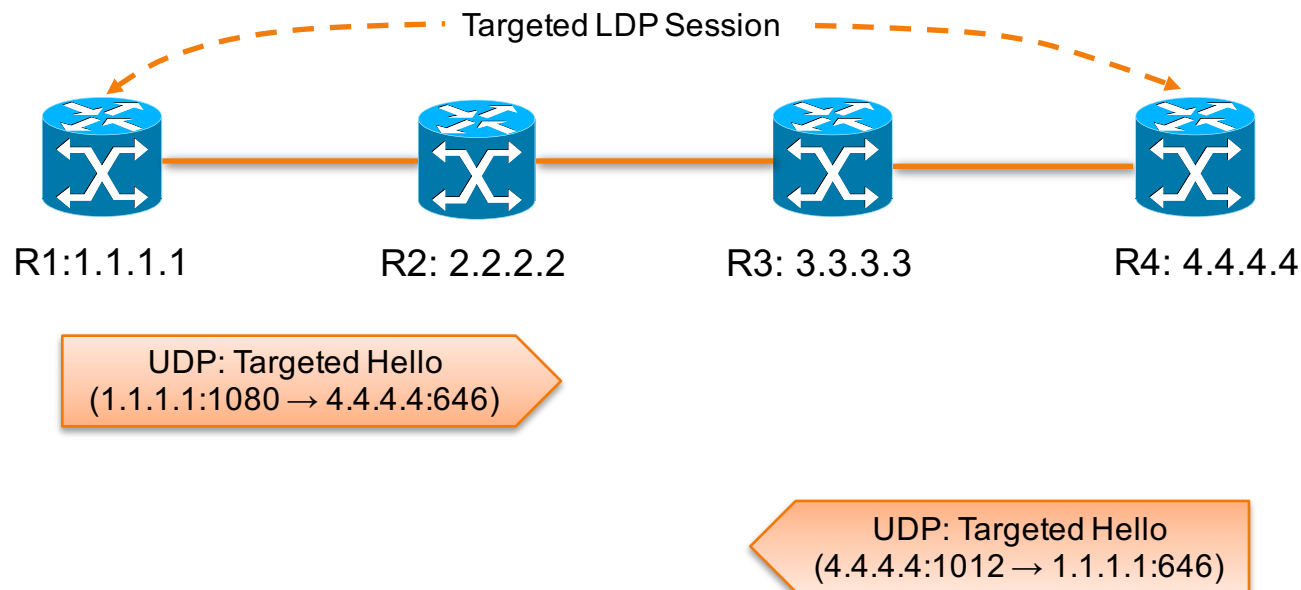
LDP Neighbor Discovery (1)

- Basic Discovery – Directly connected peer
 - LDP **Hello messages** are UDP messages that are sent on the links to the “all routers on this subnet” multicast IP address - 224.0.0.2. The UDP port used for LDP is 646.

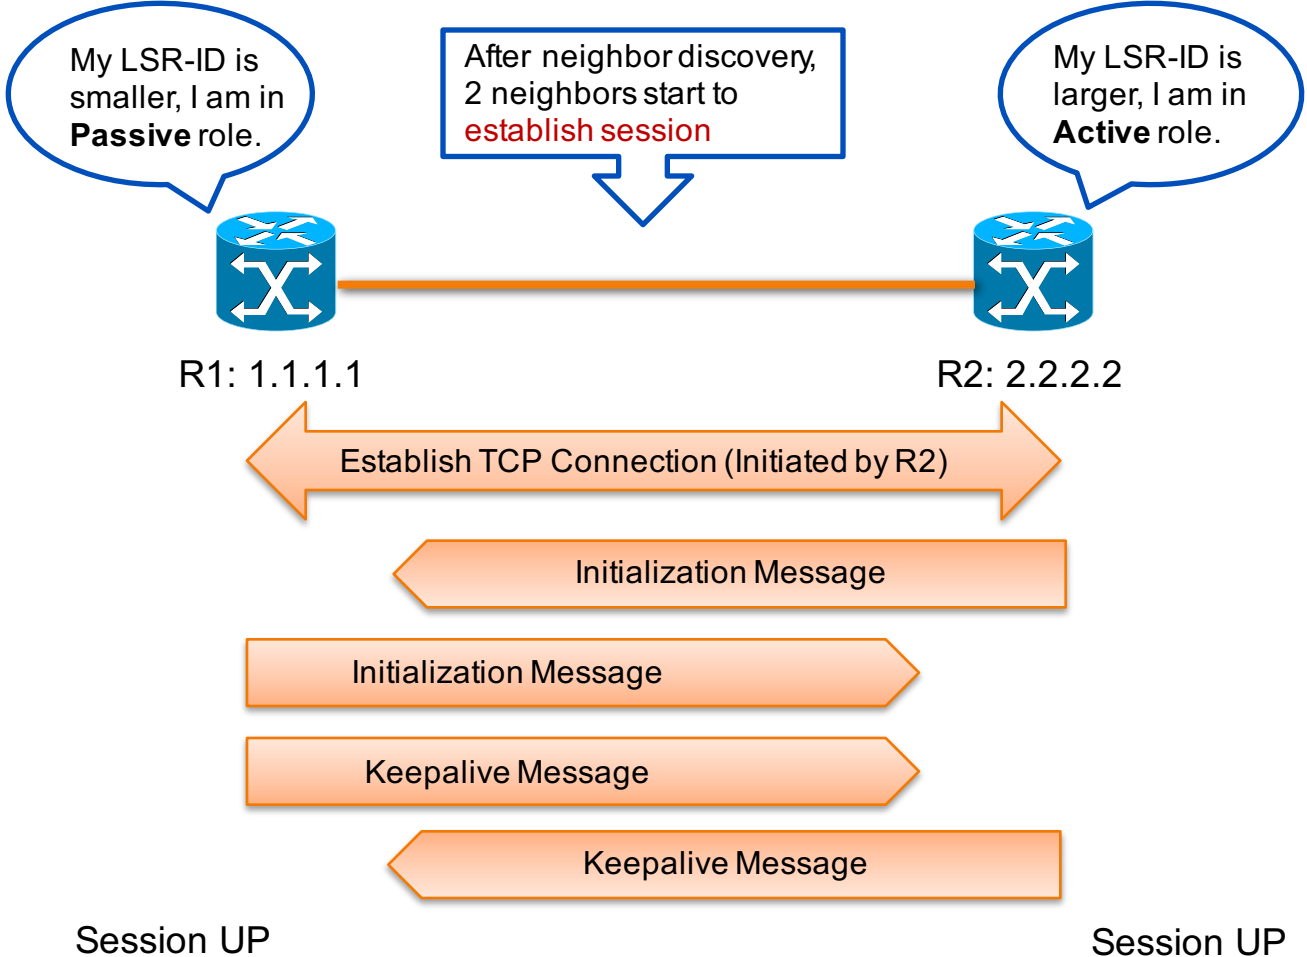


LDP Neighbor Discovery (2)

- Extended Discovery – Non-directly connected peer
 - LDP sessions between non-directly connected LSRs are supported by LDP Extended Discovery.



LDP Session Establishment and Maintenance



Label Distribution and Management

- After LDP sessions are established, labels will be distributed between LDP peers. The label distribution mode used depends on the interface and the implementation.

| | |
|--|--------------------|
| Label Distribution Control Mode | Ordered |
| | Independent |

| | |
|---------------------------------|--|
| Label Advertisement Mode | DoD (Downstream on Demand) |
| | DU (Downstream Unsolicited) |

| | |
|-----------------------------|---------------------|
| Label Retention Mode | Liberal |
| | Conservative |

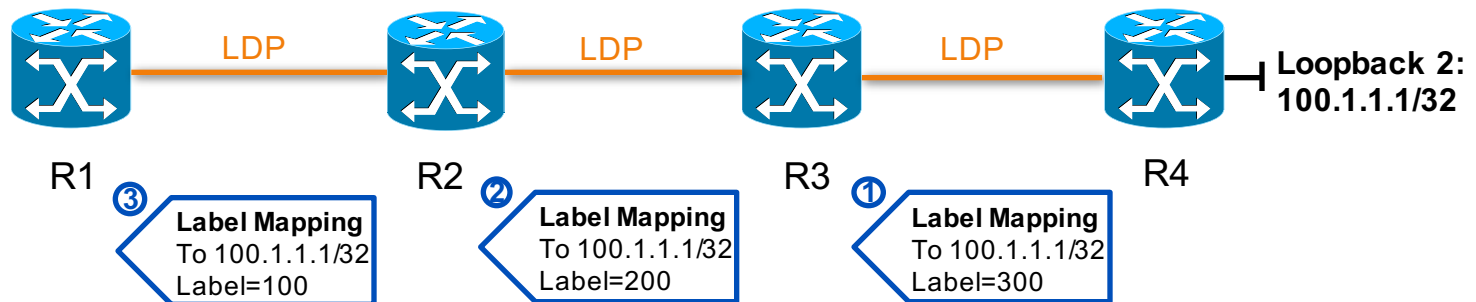
Label Distribution Control Mode

- Ordered

- In Ordered control mode, an LSR would only assign a local label for the IGP prefixes that are marked as **directly connected** in its routing table **or** also for the IGP prefixes for which it has already **received a label** from the nexthop router.

Upstream

Downstream

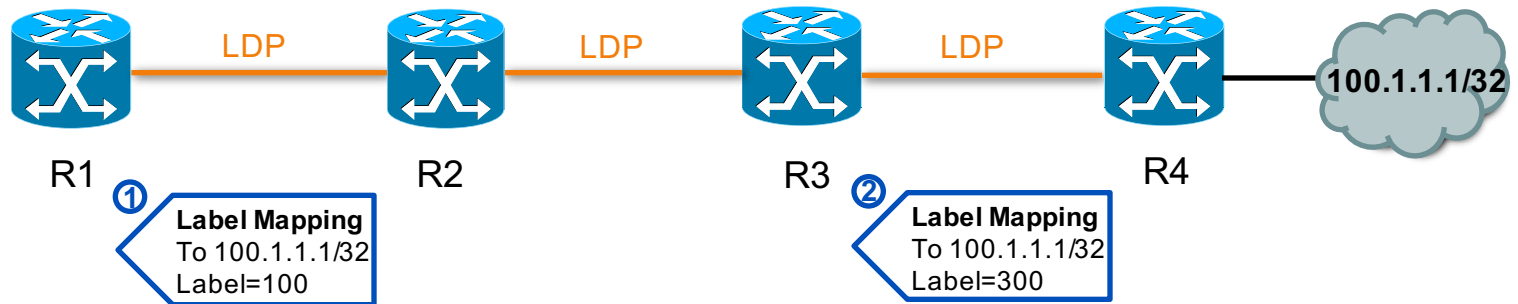


Label Distribution Control Mode - Independent

- In the independent mode, each LSR creates a local binding for a particular FEC **as soon as** it recognizes the FEC. Usually, this means that the prefix for the FEC is **in its routing table**.

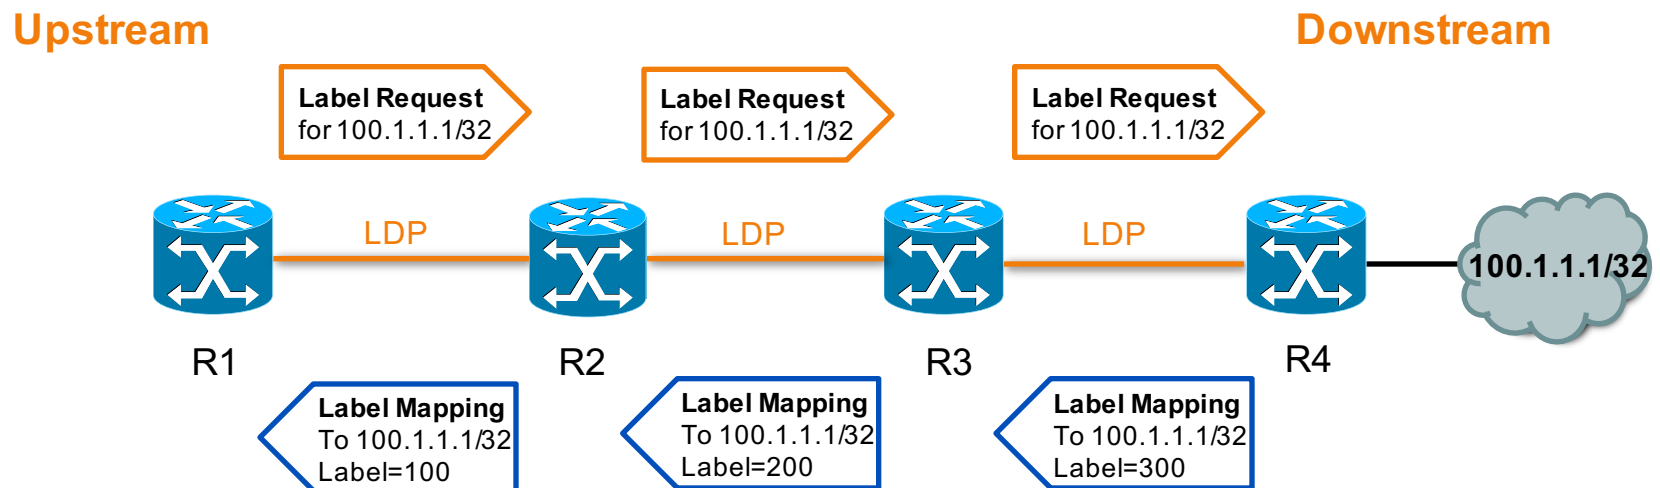
Upstream

Downstream



Label Advertisement Mode - Downstream on Demand

- In the DoD mode, an LSR distributes labels to a specified FEC only after **receiving Label Request** messages from its upstream LSR.

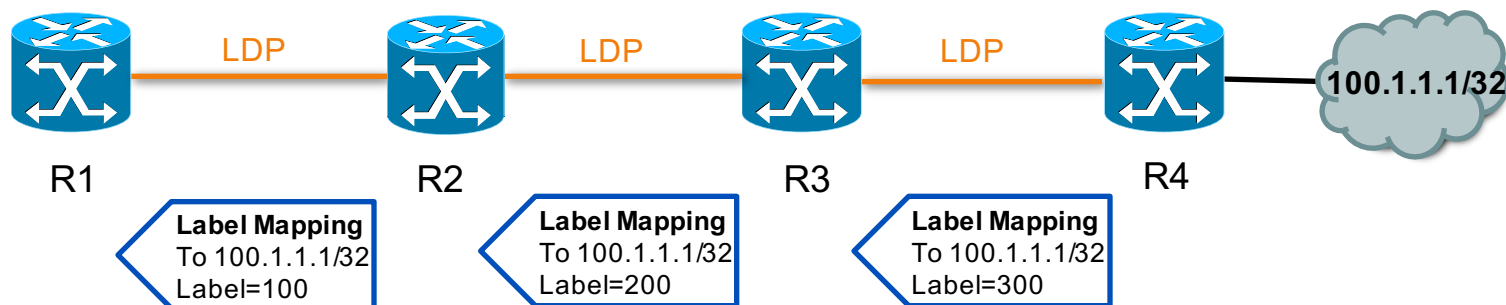


Label Advertisement Mode - Downstream Unsolicited

- In the DU mode, each LSR distributes a label to its upstream LSRs, **without** those LSRs **requesting** a label.

Upstream

Downstream

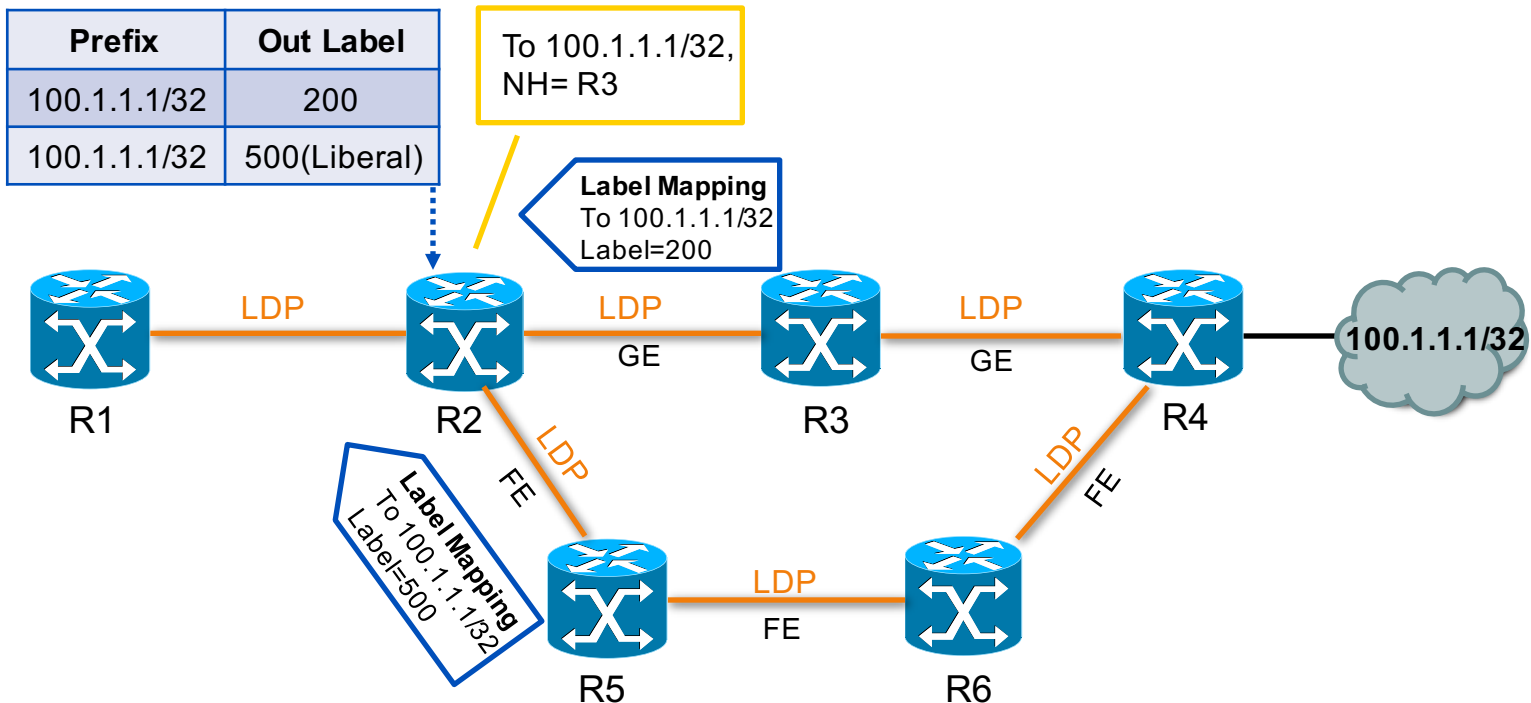


Label Retention Mode - Liberal

- In the liberal mode, an LSR **keeps all received** remote labels in the LIB, but not all are used to forward packets.

Upstream

Downstream

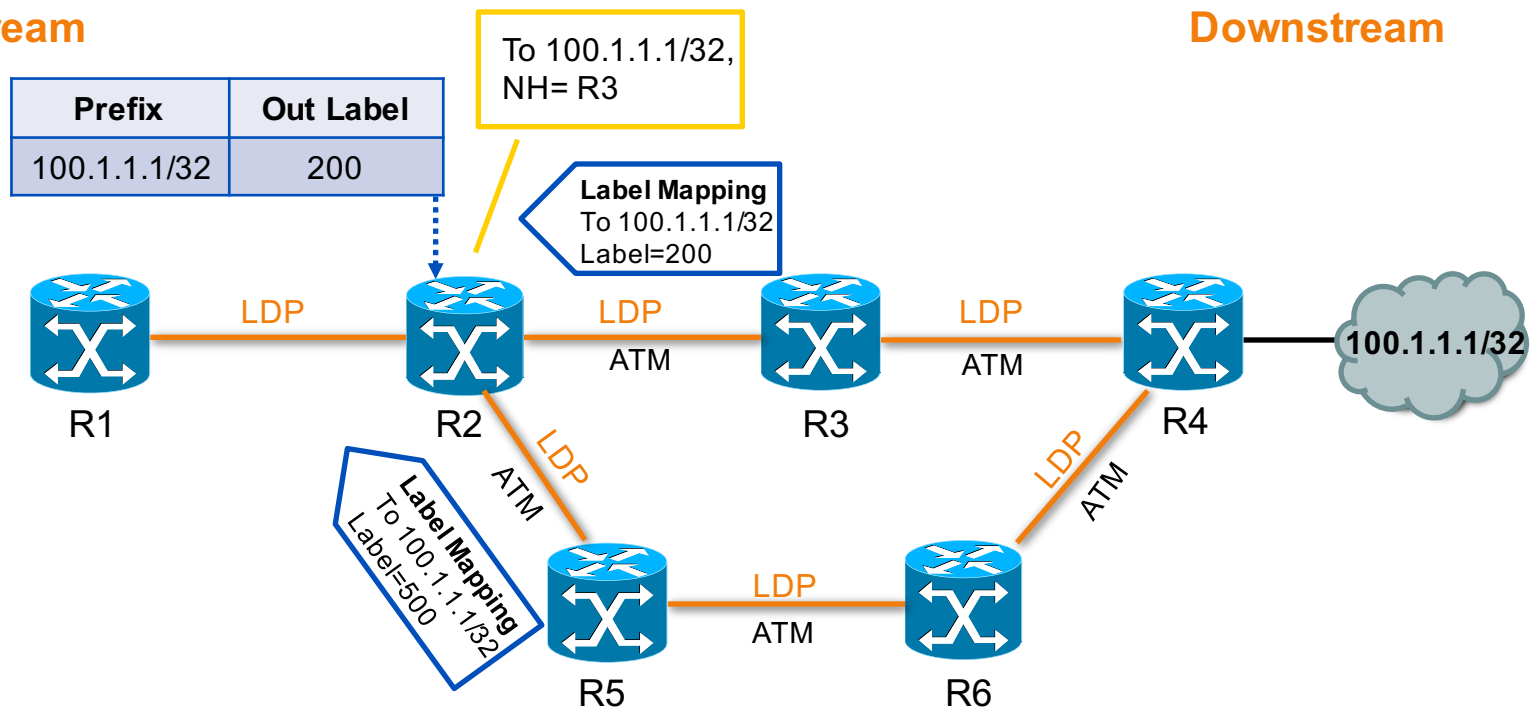


Label Retention Mode - Conservative

- An LSR that is running this mode does **not store all** remote labels in the LIB, but it stores **only** the remote label that is associated with **the next-hop** LSR for a particular FEC.

Upstream

Downstream



Label Distribution Scheme Summary

- Cisco IOS can support:

| | Control | Distribution | Retention | Label Space |
|-------------------|-------------|--------------|--------------|---------------|
| Frame Mode | Independent | DU | Liberal | Per Platform |
| Cell Mode (LCATM) | Ordered | DoD | Conservation | Per Interface |

- Junos can support:

| Control | Distribution | Retention |
|---------|--------------|-----------|
| Ordered | DU | Liberal |
| | DoD | |

- Huawei VRP can support:

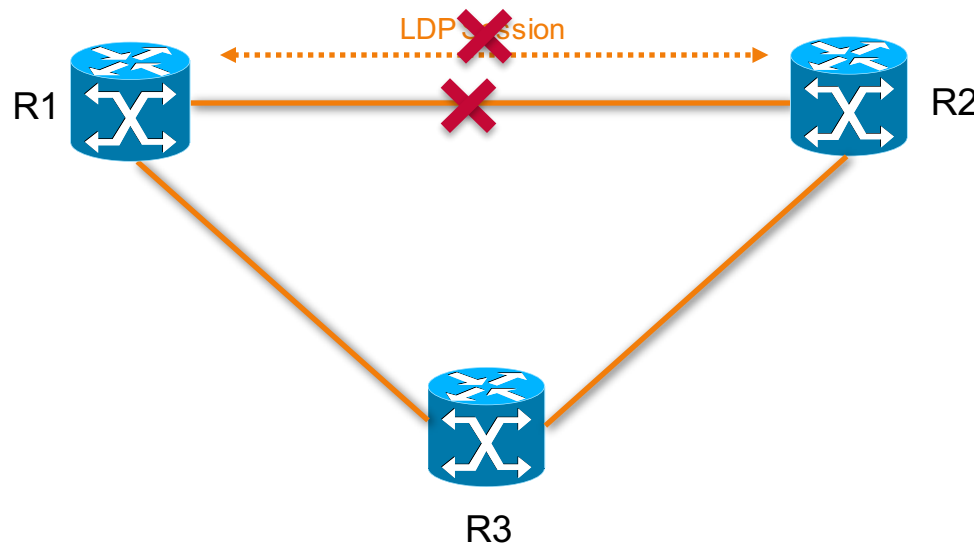
| Control | Distribution | Retention |
|---------|--------------|--------------|
| Ordered | DU | Liberal |
| Ordered | DoD | Conservation |

By default

Also support

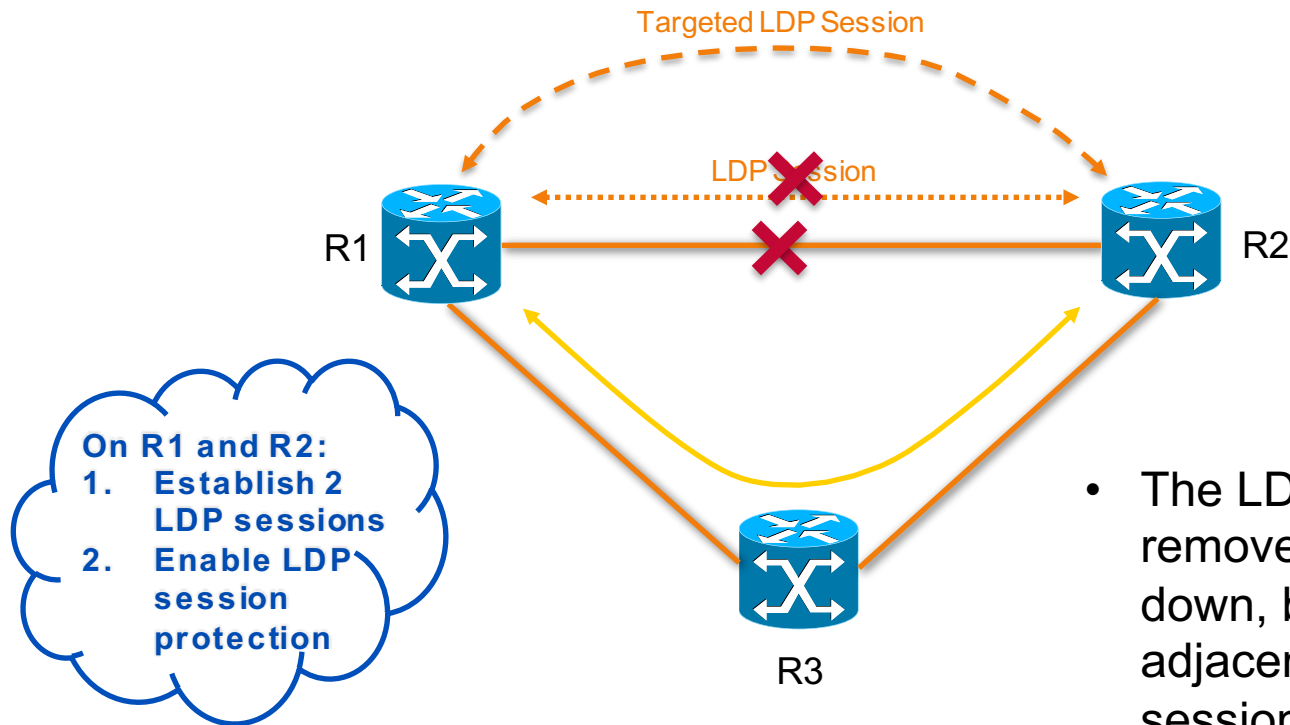
LDP Session Protection (1)

- Without LDP session protection, if the link between R1 and R2 fails, the **LDP direct link adjacency fails**.



LDP Session Protection (2)

- MPLS LDP Session Protection uses LDP Targeted Hellos to protect LDP sessions.



- The LDP link adjacency is removed when the link goes down, but the targeted adjacency keeps the LDP session up.

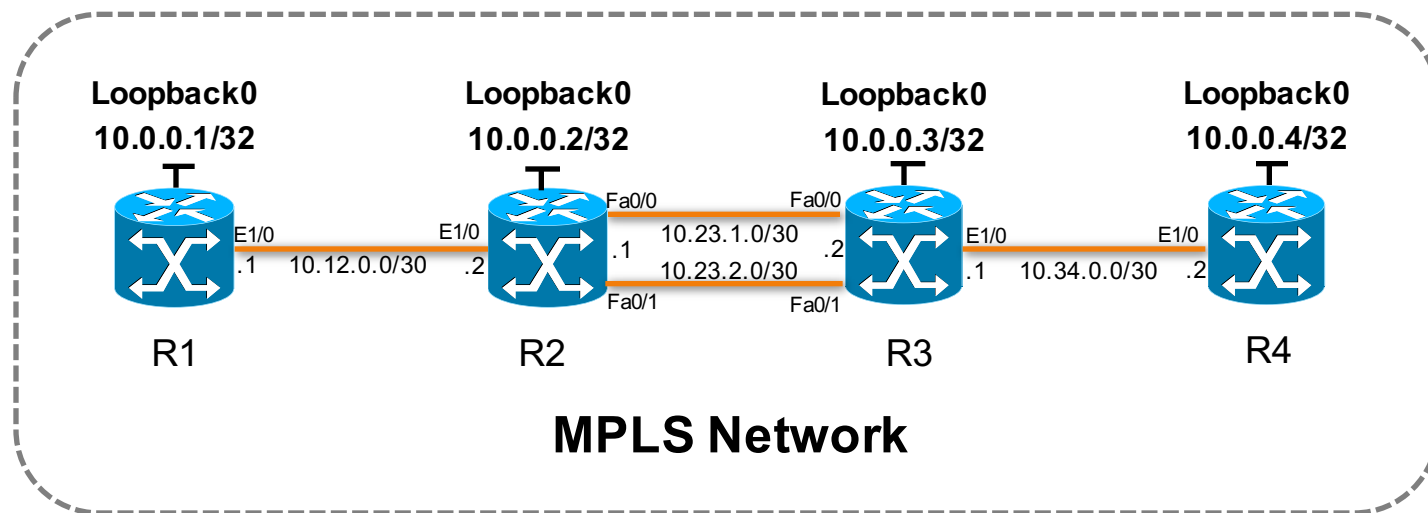
Questions?



Basic MPLS Configuration

Configuration Example

- Task: Configure MPLS LDP on Cisco IOS (Version 15.2) to set up MPLS LSP only for loopback addresses.
- Prerequisite configuration:
 - 1. IP address configuration on all the routers
 - 2. IGP configuration on all the routers



Step 1: Enable MPLS & LDP

- Configuration steps:
 - 1. Configure basic MPLS and LDP on all the routers.

R1 configuration:

```
R1 (config) # ip cef
                To make MPLS work, CEF switching is mandatory.
R1 (config) # mpls label range 100 199
                Specifying the label range for this router start from
                100 to 199.
R1 (config) # mpls ldp router-id loopback 0 force
                Forcing LDP router ID to be loopback 0 address.
R1 (config) # interface ethernet 1/0
R1 (config-if) # mpls ip
                IP MPLS is enabled on the interface.
R1 (config-if) # mpls label protocol ldp
                Label distribution protocol is LDP.
```

Check MPLS Interface

- Using **show mpls interfaces** command verifies that interfaces have been configured to use LDP:

R1#show mpls interfaces

| Interface | IP | Tunnel | BGP | Static | Operational |
|-------------|------------------|--------|-----|--------|-------------|
| Ethernet1/0 | Yes (ldp) | No | No | No | Yes |

R2#show mpls interfaces

| Interface | IP | Tunnel | BGP | Static | Operational |
|-----------------|------------------|--------|-----|--------|-------------|
| FastEthernet0/0 | Yes (ldp) | No | No | No | Yes |
| FastEthernet0/1 | Yes (ldp) | No | No | No | Yes |
| Ethernet1/0 | Yes (ldp) | No | No | No | Yes |

Check LDP Discovery

- Check LDP discovery

```
R2#show mpls ldp discovery
```

```
Local LDP Identifier:
```

```
10.0.0.2:0
```

Local LDP ID

```
Discovery Sources:
```

```
Interfaces:
```

```
FastEthernet0/0 (ldp): xmit/recv
```

```
LDP Id: 10.0.0.3:0
```

Neighbor's LDP ID

```
FastEthernet0/1 (ldp): xmit/recv
```

```
LDP Id: 10.0.0.3:0
```

```
Ethernet1/0 (ldp): xmit/recv
```

```
LDP Id: 10.0.0.1:0
```

R2 has received Hello messages from routers whose ID are 10.0.0.3 and 10.0.0.1

Check LDP Neighbors

- Check LDP neighbors on R1

```
R1#show mpls ldp neighbor
```

```
Peer LDP Ident: 10.0.0.2:0; Local LDP Ident 10.0.0.1:0
```

```
TCP connection: 10.0.0.2.48548 - 10.0.0.1.646
```

```
State: Oper; Msgs sent/rcvd: 34/34; Downstream
```

```
Up time: 00:09:57
```

```
LDP discovery sources:
```

```
 Ethernet1/0, Src IP addr: 10.12.0.2
```

```
Addresses bound to peer LDP Ident:
```

```
10.23.1.1
```

```
10.23.2.1
```

```
10.12.0.2
```

```
10.0.0.2
```

LDP session is
a TCP session
(port = 646)

Operational is the
stable state of LDP
session.

Check LDP Neighbors

- Check LDP neighbors on R2

```
R2#show mpls ldp neighbor
```

```
Peer LDP Ident: 10.0.0.3:0; Local LDP Ident 10.0.0.2:0
```

```
TCP connection: 10.0.0.3.28664 - 10.0.0.2.646
```

```
State: Oper; Msgs sent/rcvd: 36/36; Downstream
```

```
Up time: 00:12:12
```

```
LDP discovery sources:
```

```
FastEthernet0/0, Src IP addr: 10.23.1.2
```

```
FastEthernet0/1, Src IP addr: 10.23.2.2
```

```
Addresses bound to peer LDP Ident:
```

```
10.23.1.2      10.23.2.2      10.34.0.1      10.0.0.3
```

```
Peer LDP Ident: 10.0.0.1:0; Local LDP Ident 10.0.0.2:0
```

```
TCP connection: 10.0.0.1.646 - 10.0.0.2.48548
```

```
State: Oper; Msgs sent/rcvd: 36/36; Downstream
```

```
Up time: 00:11:40
```


```
LDP discovery sources:
```

```
Ethernet1/0, Src IP addr: 10.12.0.1
```

```
Addresses bound to peer LDP Ident:
```

```
10.12.0.1      10.0.0.1
```

Multiple links
between two routers
still mean single
LDP session.



Check LDP Label Information Base

- Check LDP LIB

```
R2#show mpls ldp bindings
```

```
lib entry: 10.0.0.1/32, rev 24
  local binding: label: 206
  remote binding: lsr: 10.0.0.3:0, label: 302
  remote binding: lsr: 10.0.0.1:0, label: imp-null
lib entry: 10.0.0.2/32, rev 18
  local binding: label: imp-null
  remote binding: lsr: 10.0.0.3:0, label: 301
  remote binding: lsr: 10.0.0.1:0, label: 101
lib entry: 10.0.0.3/32, rev 22
  local binding: label: 205
  remote binding: lsr: 10.0.0.3:0, label: imp-null
  remote binding: lsr: 10.0.0.1:0, label: 100
  .....(omitted)
```

```
R2#show mpls ldp bindings 10.0.0.4 32
```

```
lib entry: 10.0.0.4/32, rev 20
  local binding: label: 204
  remote binding: lsr: 10.0.0.3:0, label: 300
  remote binding: lsr: 10.0.0.1:0, label: 107
```

Check Label Forwarding Table

- Check label forwarding table

```
R2#show mpls forwarding-table
```

| Local Label | Outgoing Label | Prefix or Tunnel Id | Bytes Switched | Label | Outgoing interface | Next Hop |
|-------------|----------------|---------------------|----------------|-------|--------------------|-----------|
| 203 | Pop Label | 10.34.0.0/30 | 0 | | Fa0/0 | 10.23.1.2 |
| | Pop Label | 10.34.0.0/30 | 0 | | Fa0/1 | 10.23.2.2 |
| 204 | 300 | 10.0.0.4/32 | 0 | | Fa0/0 | 10.23.1.2 |
| | 300 | 10.0.0.4/32 | 0 | | Fa0/1 | 10.23.2.2 |
| 205 | Pop Label | 10.0.0.3/32 | 0 | | Fa0/0 | 10.23.1.2 |
| | Pop Label | 10.0.0.3/32 | 0 | | Fa0/1 | 10.23.2.2 |
| 206 | Pop Label | 10.0.0.1/32 | 0 | | Et1/0 | 10.12.0.1 |

Step 2: Configure Conditional Label Distribution

- Only set up LSP for loopback addresses.
 - 2.1 Create prefix-list on each router, all the loopback addresses are in 10.0.0.0/24 block.

```
R1(config)#ip prefix-list ALL-LOOPBACK seq 5 permit  
10.0.0.0/24 le 32
```

- 2.2 Apply the prefix-list

```
R1(config)#mpls ldp label  
R1(config-ldp-lbl)#allocate global prefix-list ALL-LOOPBACK
```

Allocate labels for the routes matching ALL-LOOPBACK prefix-list.

Verify the Results of Conditional Label Distribution

- Before the configuration.

```
R1#show mpls ldp bindings
```

```
lib entry: 10.0.0.1/32, rev 51
    local binding: label: imp-null
    remote binding: lsr: 10.0.0.2:0, label: 206
lib entry: 10.0.0.2/32, rev 52
    local binding: label: 101
    remote binding: lsr: 10.0.0.2:0, label: imp-null
lib entry: 10.0.0.3/32, rev 53
    local binding: label: 100
    remote binding: lsr: 10.0.0.2:0, label: 205
lib entry: 10.0.0.4/32, rev 54
    local binding: label: 107
    remote binding: lsr: 10.0.0.2:0, label: 204
lib entry: 10.12.0.0/30, rev 71
    local binding: label: imp-null
    remote binding: lsr: 10.0.0.2:0, label: imp-null
lib entry: 10.23.1.0/30, rev 72
    local binding: label: 102
    remote binding: lsr: 10.0.0.2:0, label: imp-null
lib entry: 10.23.2.0/30, rev 73
    local binding: label: 103
    remote binding: lsr: 10.0.0.2:0, label: imp-null
lib entry: 10.34.0.0/30, rev 75
    local binding: label: 104
    remote binding: lsr: 10.0.0.2:0, label: 200
```

Entries for **all** the prefixes in IP routing table.

Verify the Results of Conditional Label Distribution

- After configure on all the routers.

R1#show mpls ldp bindings

```
lib entry: 10.0.0.1/32, rev 51
    local binding: label: imp-null
    remote binding: lsr: 10.0.0.2:0, label: 206
lib entry: 10.0.0.2/32, rev 52
    local binding: label: 101
    remote binding: lsr: 10.0.0.2:0, label: imp-null
lib entry: 10.0.0.3/32, rev 53
    local binding: label: 100
    remote binding: lsr: 10.0.0.2:0, label: 205
lib entry: 10.0.0.4/32, rev 54
    local binding: label: 107
    remote binding: lsr: 10.0.0.2:0, label: 204
```

Only the entries for loopback addresses.

Check the LSP

- Check the LSP for 10.0.0.4 from R1 to R3

```
R1#show mpls forwarding-table 10.0.0.4 32
Local      Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label      Label     or Tunnel Id   Switched     interface
107        204      10.0.0.4/32    0            Et1/0     10.12.0.2
```

```
R2#show mpls forwarding-table 10.0.0.4 32
Local      Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label      Label     or Tunnel Id   Switched     interface
204        300      10.0.0.4/32    0            Fa0/0     10.23.1.2
          300      10.0.0.4/32    0            Fa0/1     10.23.2.2
```

```
R3# show mpls forwarding-table 10.0.0.4 32
Local      Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label      Label     or Tunnel Id   Switched     interface
300        Pop Label 10.0.0.4/32    0            Fa0/0     10.34.0.2
```

Implicit Null

If I want to use an outgoing label at the penultimate hop for keeping QoS info. What can I do?

Explicit-null Label

- Explicit Null label can be used to keep the QoS information.

Explicit Null (IPv4) = Label 0

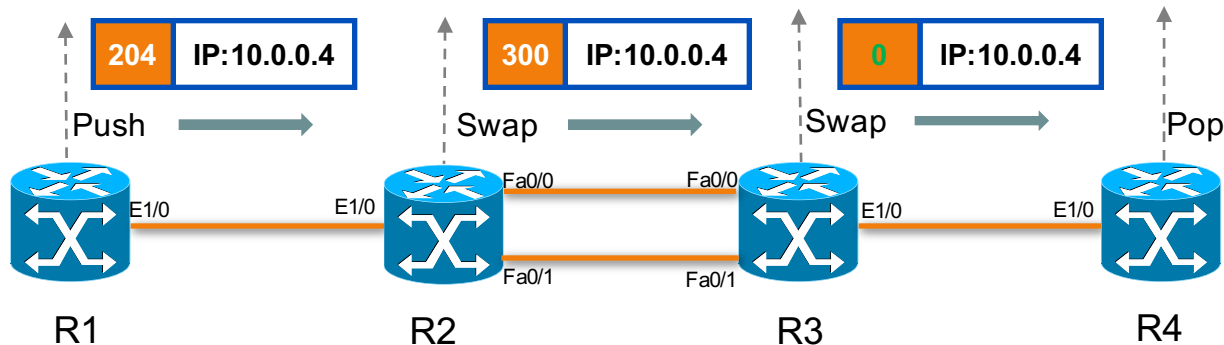
Explicit Null (IPv6) = Label 2

| | |
|---------------------|------|
| Prefix: 10.0.0.4/32 | |
| Local Label | Null |
| Out Interface | E1/0 |
| Out Label | 204 |
| Operation | Push |

| | |
|---------------------|-------------|
| Prefix: 10.0.0.4/32 | |
| Local Label | 204 |
| Out Interface | Fa0/0 Fa0/1 |
| Out Label | 300 |
| Operation | Swap |

| | |
|---------------------|------------|
| Prefix: 10.0.0.4/32 | |
| Local Label | 300 |
| Out Interface | E1/0 |
| Out Label | Explicit-n |
| Operation | Swap |

| | |
|---------------------|------------|
| Prefix: 10.0.0.4/32 | |
| Local Label | Explicit-n |
| Out Interface | -- |
| Out Label | -- |
| Operation | POP |



Additional Task: Using Explicit-null Label

- Explicit-null configuration on R4:

```
R4(config)# mpls ldp explicit-null
```

- After configuring this command, check the label forwarding table on R3.

```
R3#show mpls forwarding-table 10.0.0.4 32
Local      Outgoing  Prefix          Bytes Label  Outgoing  Next Hop
Label     Label     or Tunnel Id   Switched     interface
300      explicit-n 10.0.0.4/32    0            Et1/0     10.34.0.2
```

Out label is
explicit-null

Questions?

