

# SDN Workshop

Contact: [training@apnic.net](mailto:training@apnic.net)



# Segment Routing - Lab

SDN Workshop

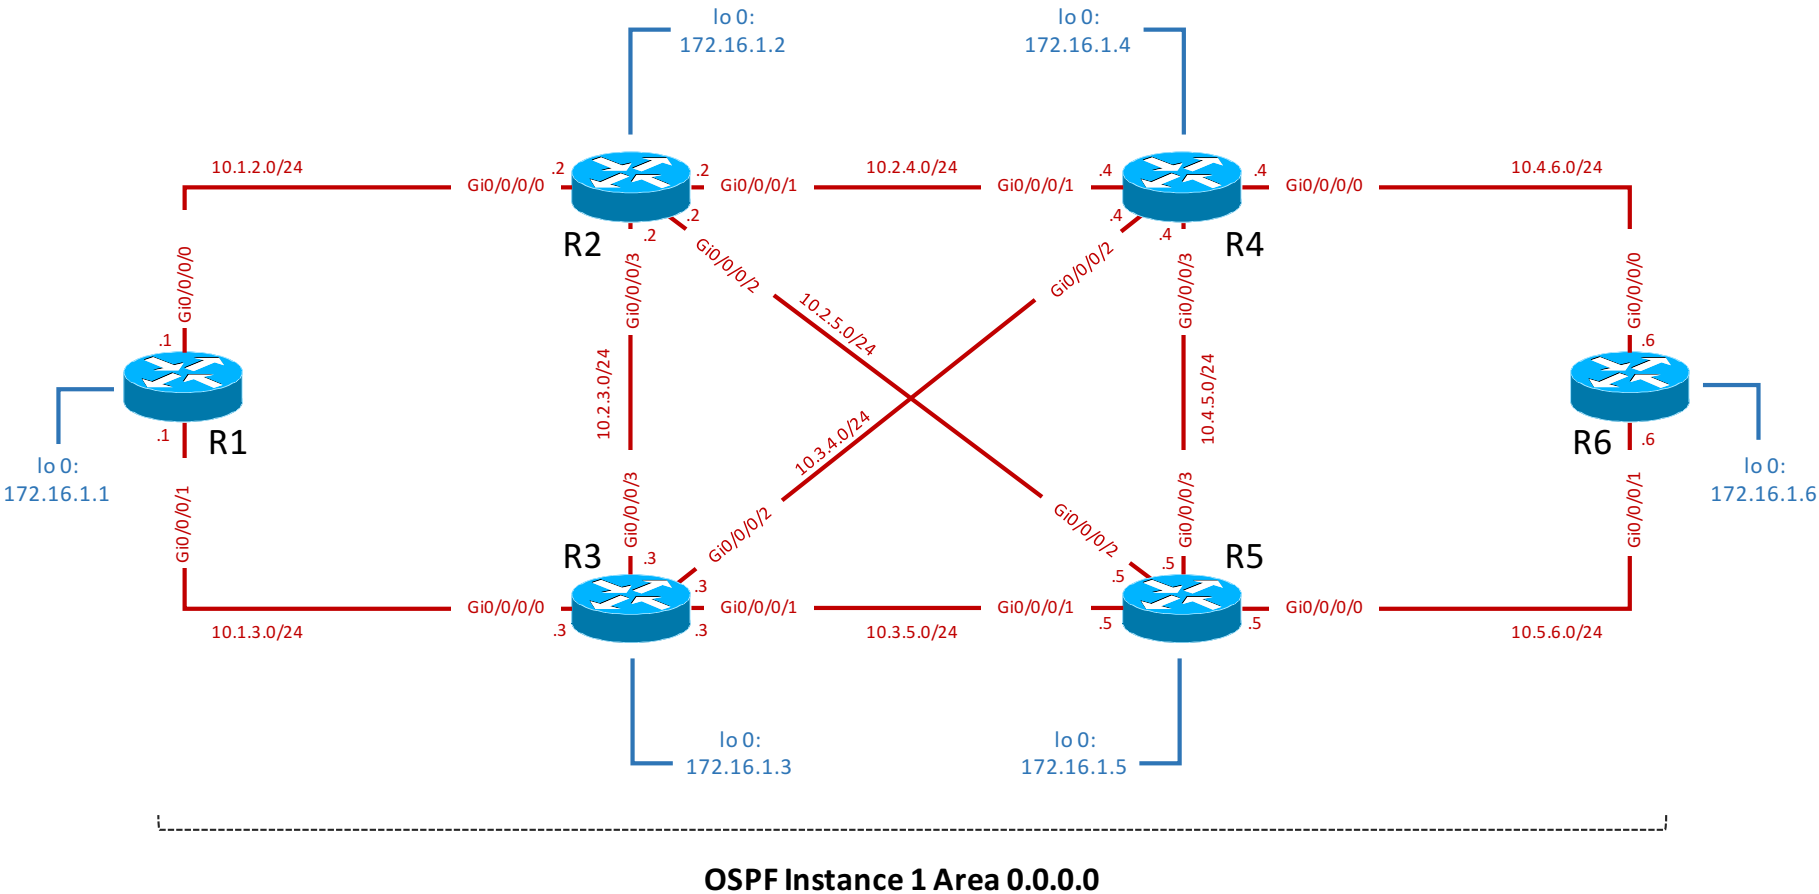


# Segment Routing Configuration

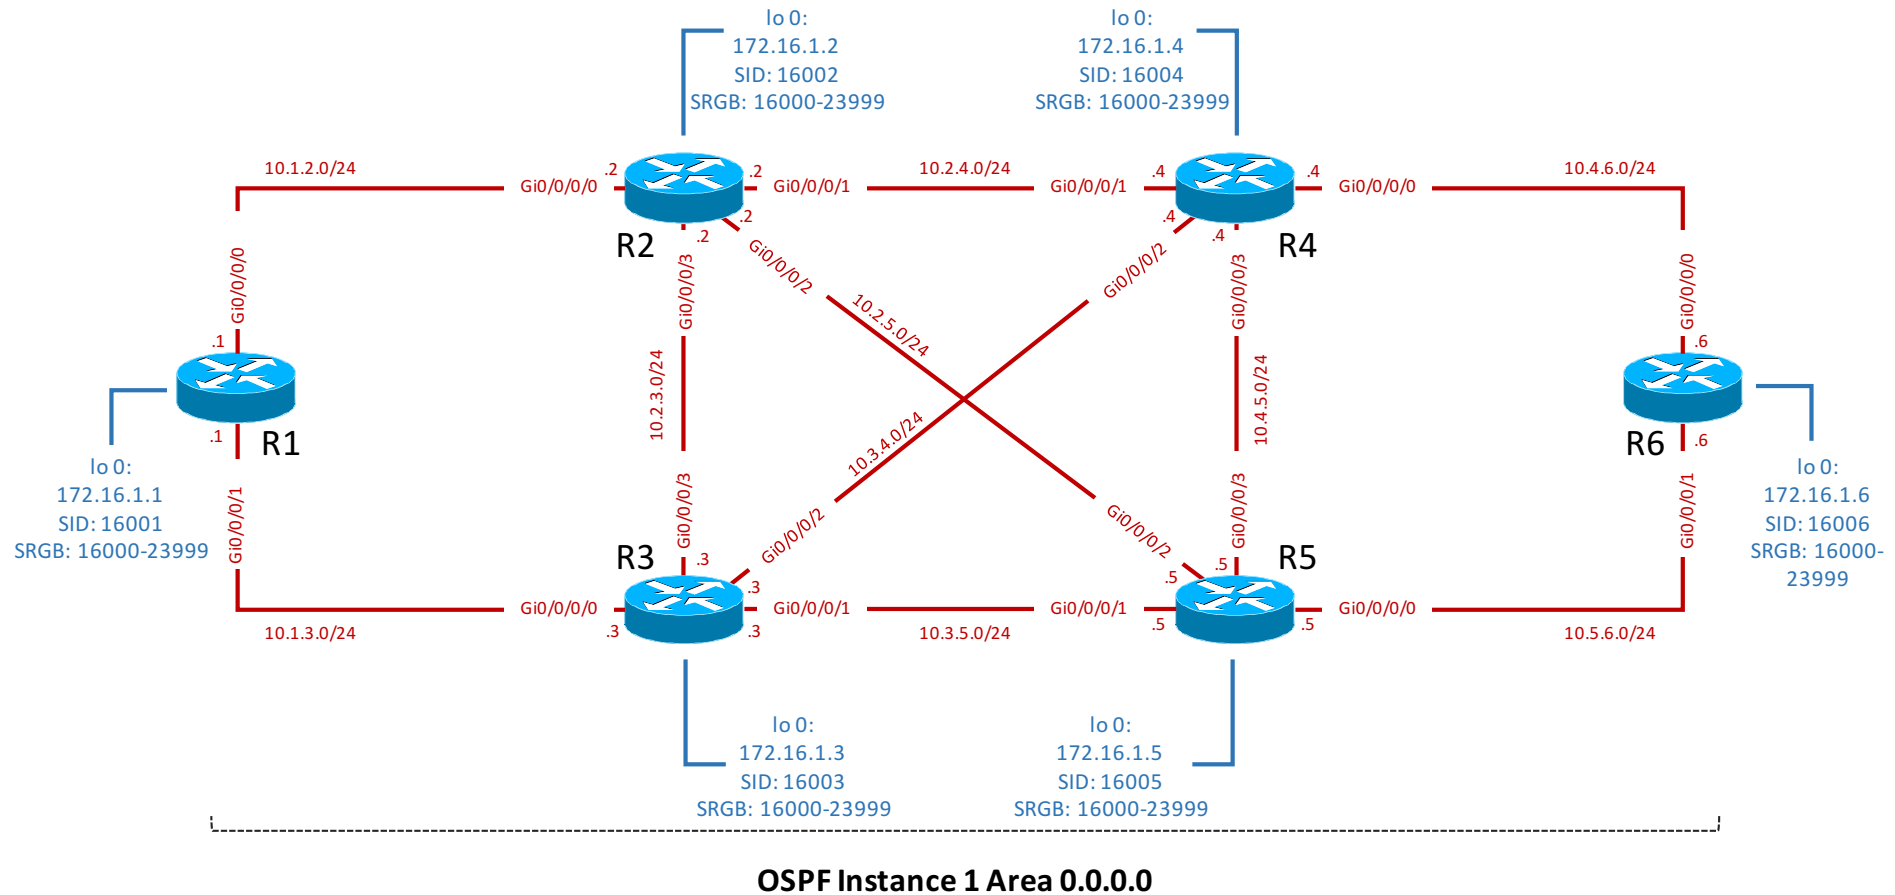
# Configuration Example

- Task: configure Segment Routing on Cisco IOS-XRv 6.1.2 nodes with various parameter settings
- Pre-requisite configuration:
  - IP address configuration on all the routers
  - IGP (OSPF) configuration on all the routers

# Base Topology



# Part 1: Basic Segment Routing



# Enable basic Segment Routing

- Enable MPLS-TE and segment routing, configure SRGB and prefix SID for loopback interface
- R1 configuration:

```
R1(config)# mpls traffic-eng
                ! Enable traffic engineering functionality on the node.
R1(config)# router ospf 1
R1(config-ospf)# segment-routing global-block 16000 23999
                ! Configure the SRGB for this node
R1(config-ospf)# segment-routing mpls
                ! Enables Segment Routing using the mpls dataplane
R1(config-ospf)# mpls traffic-eng router-id loopback0
                ! Configures loopback 0 as the traffic-engineering router identifier.
R1(config-ospf)# area 0.0.0.0
R1(config-ospf-ar)# mpls traffic-eng
                ! Enables traffic-engineering functionality for OSPF.
R1(config-ospf-ar)# interface loopback 0
R1(config-ospf-ar)# prefix-sid absolute 16001
                ! Assigns an absolute SID value to loopback 0
```

# Check SRGB

## – Check SRGB

```
RP/0/0/CPU0:R1#show mpls label range
Range for dynamic labels: Min/Max: 24000/1048575
```

```
RP/0/0/CPU0:R1#show mpls label table detail
```

Table	Label	Owner	State	Rewrite
0	0	LSD(A)	InUse	Yes
0	1	LSD(A)	InUse	Yes
0	2	LSD(A)	InUse	Yes
0	13	LSD(A)	InUse	Yes
0	16000	OSPF(A):ospf-1 (Lbl-blk SRGB, vers:0, (start label=16000, size=8000))	InUse	No
0	24000	OSPF(A):ospf-1 (SR Adj Segment IPv4, vers:0, index=0, type=1, intf=Gi0/0/0/0, nh=10.1.2.2)	InUse	Yes
0	24001	OSPF(A):ospf-1 (SR Adj Segment IPv4, vers:0, index=0, type=2, intf=Gi0/0/0/0, nh=10.1.2.2)	InUse	Yes
0	24002	OSPF(A):ospf-1 (SR Adj Segment IPv4, vers:0, index=0, type=1, intf=Gi0/0/0/1, nh=10.1.3.3)	InUse	Yes
0	24003	OSPF(A):ospf-1 (SR Adj Segment IPv4, vers:0, index=0, type=2, intf=Gi0/0/0/1, nh=10.1.3.3)	InUse	Yes

SRGB



# Check label table

- Check label table

```
RP/0/0/CPU0:R1#show mpls forwarding
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16002	Pop	SR Pfx (idx 2)	Gi0/0/0/0	10.1.2.2	0
16003	Pop	SR Pfx (idx 3)	Gi0/0/0/1	10.1.3.3	0
16004	16004	SR Pfx (idx 4)	Gi0/0/0/0	10.1.2.2	0
	16004	SR Pfx (idx 4)	Gi0/0/0/1	10.1.3.3	0
16005	16005	SR Pfx (idx 5)	Gi0/0/0/0	10.1.2.2	0
	16005	SR Pfx (idx 5)	Gi0/0/0/1	10.1.3.3	0
16006	16006	SR Pfx (idx 6)	Gi0/0/0/0	10.1.2.2	0
	16006	SR Pfx (idx 6)	Gi0/0/0/1	10.1.3.3	0
24000	Pop	SR Adj (idx 0)	Gi0/0/0/0	10.1.2.2	0
24001	Pop	SR Adj (idx 0)	Gi0/0/0/0	10.1.2.2	0
24002	Pop	SR Adj (idx 0)	Gi0/0/0/1	10.1.3.3	0
24003	Pop	SR Adj (idx 0)	Gi0/0/0/1	10.1.3.3	0

Prefix SIDs

Adjacency SIDs

# Check OSPF opaque LSAs (1)

## – Check OSPF opaque LSAs

```
RP/0/0/CPU0:R1#show ospf database opaque-area self-originate
```

```
LS age: 1696
```

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

```
LS Type: Opaque Area Link
```

```
Link State ID: 4.0.0.0
```

```
Opaque Type: 4
```

```
Opaque ID: 0
```

```
Advertising Router: 172.16.1.1
```

```
LS Seq Number: 80000001
```

```
Checksum: 0xaa0b
```

```
Length: 52
```

```
Router Information TLV: Length: 4
```

```
Capabilities:
```

```
Graceful Restart Helper Capable
```

```
Stub Router Capable
```

```
All capability bits: 0x60000000
```

```
Segment Routing Algorithm TLV: Length: 2
```

```
Algorithm: 0
```

```
Algorithm: 1
```

```
Segment Routing Range TLV: Length: 12
```

```
Range Size: 8000
```

```
SID sub-TLV: Length 3
```

```
Label: 16000
```

# Check OSPF opaque LSAs (2)

## – Check OSPF opaque LSAs

```
RP/0/0/CPU0:R1#show ospf database opaque-area self-originate
```

```
LS age: 1614
```

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

```
LS Type: Opaque Area Link
```

```
Link State ID: 7.0.0.1
```

```
Opaque Type: 7
```

```
Opaque ID: 1
```

```
Advertising Router: 172.16.1.1
```

```
LS Seq Number: 80000001
```

```
Checksum: 0x79a7
```

```
Length: 44
```

```
Extended Prefix TLV: Length: 20
```

```
Route-type: 1
```

```
AF      : 0
```

```
Flags   : 0x40
```

```
Prefix  : 172.16.1.1/32
```

```
SID sub-TLV: Length: 8
```

```
Flags   : 0x0
```

```
MTID    : 0
```

```
Algo    : 0
```

```
SID Index : 1
```

# Check CEF entry

- Check CEF entry

```
RP/0/0/CPU0:R1#show ip cef 172.16.1.6/32
172.16.1.6/32, version 46, internal 0x1000001 0x81 (ptr 0xa1421ff4) [1], 0x0
(0xa13ed5f0), 0xa28 (0xa17720a8)
Updated Aug 25 10:46:19.328
local adjacency 10.1.2.2
Prefix Len 32, traffic index 0, precedence n/a, priority 1
via 10.1.2.2/32, GigabitEthernet0/0/0/0, 7 dependencies, weight 0, class 0
[flags 0x0]
  path-idx 0 NHID 0x0 [0xa10fe250 0x0]
  next hop 10.1.2.2/32
  local adjacency
    local label 16006      labels imposed {16006}
  via 10.1.3.3/32, GigabitEthernet0/0/0/1, 7 dependencies, weight 0, class 0
[flags 0x0]
  path-idx 1 NHID 0x0 [0xa10fe2f8 0x0]
  next hop 10.1.3.3/32
  local adjacency
    local label 16006      labels imposed {16006}
```

ECMP

# Check details for R6 prefix SID label

- Check details for R6 prefix SID label

```
RP/0/0/CPU0:R1#show mpls forwarding labels 16006 detai
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
-------------	----------------	--------------	--------------------	----------	----------------

16006	16006	SR Pfx (idx 6)	Gi0/0/0/0	10.1.2.2	0
-------	-------	----------------	-----------	----------	---

Updated: Aug 25 10:46:19.328

Version: 46, Priority: 1

Label Stack (Top -> Bottom): { 16006 }

NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0

MAC/Encaps: 14/18, MTU: 1500

Packets Switched: 0

16006		SR Pfx (idx 6)	Gi0/0/0/1	10.1.3.3	0
-------	--	----------------	-----------	----------	---

Updated: Aug 25 10:46:19.328

Version: 46, Priority: 1

Label Stack (Top -> Bottom): { 16006 }

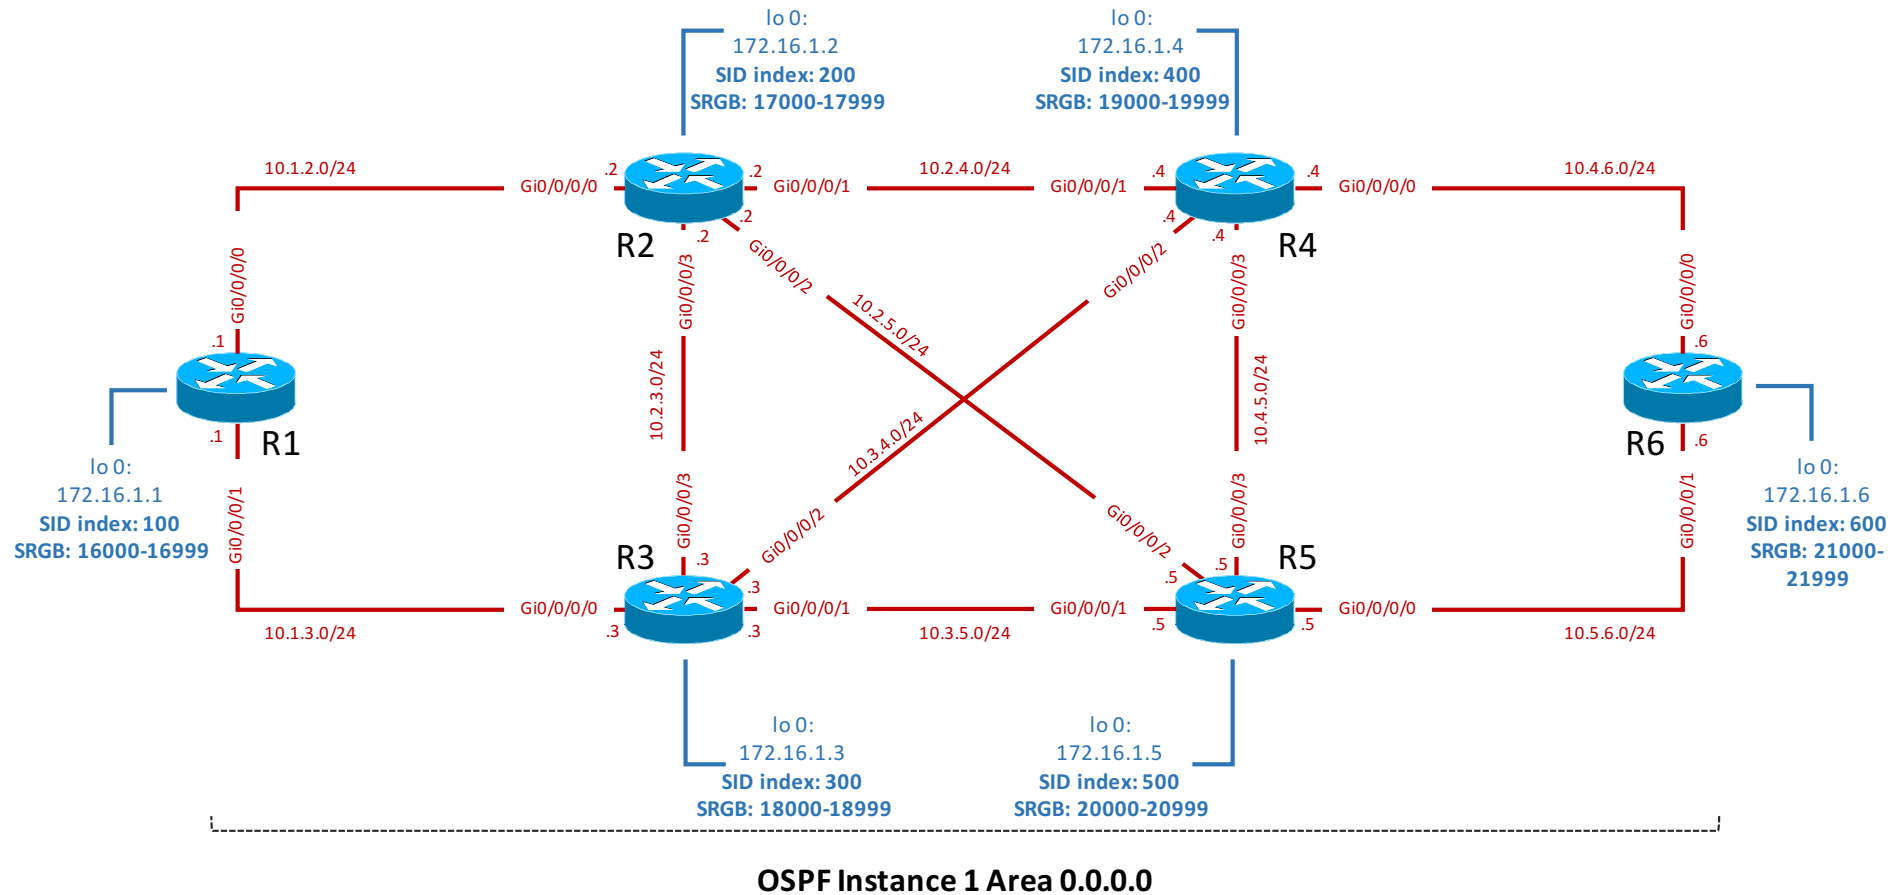
NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0

MAC/Encaps: 14/18, MTU: 1500

Packets Switched: 0

Traffic-Matrix Packets/Bytes Switched: 0/0

# Part 2: Different SRGBs



# Configure SID indexes and SRGBs

- Assign index SID to loopback 0

```
R1(config)# router ospf 1
R1(config-ospf)# area 0.0.0.0
R1(config-ospf-ar)# interface loopback 0
R1(config-ospf-ar)# prefix-sid index 100
```

- Allocate a new router-specific SRGB

```
R1(config)# router ospf 1
R1(config-ospf)# segment-routing global-block 16000 16999
```

# Check SRGB

## – Check SRGB

```
RP/0/0/CPU0:R1#show mpls label range
Range for dynamic labels: Min/Max: 24000/1048575
```

```
RP/0/0/CPU0:R1#show mpls label table detail
Table Label  Owner                               State Rewrite
-----
0      0      LSD(A)                               InUse  Yes
0      1      LSD(A)                               InUse  Yes
0      2      LSD(A)                               InUse  Yes
0      13     LSD(A)                               InUse  Yes
0      16000  OSPF(A):ospf-1                       InUse  No
(Lbl-blk SRGB, vers:0, (start label=16000, size=1000))
0      24000  OSPF(A):ospf-1                       InUse  Yes
(SR Adj Segment IPv4, vers:0, index=0, type=1, intf=Gi0/0/0/0, nh=10.1.2.2)
0      24001  OSPF(A):ospf-1                       InUse  Yes
(SR Adj Segment IPv4, vers:0, index=0, type=2, intf=Gi0/0/0/0, nh=10.1.2.2)
0      24002  OSPF(A):ospf-1                       InUse  Yes
(SR Adj Segment IPv4, vers:0, index=0, type=1, intf=Gi0/0/0/1, nh=10.1.3.3)
0      24003  OSPF(A):ospf-1                       InUse  Yes
(SR Adj Segment IPv4, vers:0, index=0, type=2, intf=Gi0/0/0/1, nh=10.1.3.3)
```

SRGB



# Check label table

- Check label table

```
RP/0/0/CPU0:R1#show mpls forwarding
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
16200	Pop	SR Pfx (idx 200)	Gi0/0/0/0	10.1.2.2	0
16300	Pop	SR Pfx (idx 300)	Gi0/0/0/1	10.1.3.3	0
16400	17400	SR Pfx (idx 400)	Gi0/0/0/0	10.1.2.2	0
	18400	SR Pfx (idx 400)	Gi0/0/0/1	10.1.3.3	0
16500	17500	SR Pfx (idx 500)	Gi0/0/0/0	10.1.2.2	0
	18500	SR Pfx (idx 500)	Gi0/0/0/1	10.1.3.3	0
16600	17600	SR Pfx (idx 600)	Gi0/0/0/0	10.1.2.2	0
	18600	SR Pfx (idx 600)	Gi0/0/0/1	10.1.3.3	0
24000	Pop	SR Adj (idx 0)	Gi0/0/0/0	10.1.2.2	0
24001	Pop	SR Adj (idx 0)	Gi0/0/0/0	10.1.2.2	0
24002	Pop	SR Adj (idx 0)	Gi0/0/0/1	10.1.3.3	0
24003	Pop	SR Adj (idx 0)	Gi0/0/0/1	10.1.3.3	0

Prefix SIDs

Adjacency SIDs

# Check OSPF opaque LSAs (1)

## – Check OSPF opaque LSAs

```
RP/0/0/CPU0:R1#show ospf database opaque-area self-originate
```

```
LS age: 454
```

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

```
LS Type: Opaque Area Link
```

```
Link State ID: 4.0.0.0
```

```
Opaque Type: 4
```

```
Opaque ID: 0
```

```
Advertising Router: 172.16.1.1
```

```
LS Seq Number: 80000001
```

```
Checksum: 0xf028
```

```
Length: 52
```

```
Router Information TLV: Length: 4
```

```
Capabilities:
```

```
Graceful Restart Helper Capable
```

```
Stub Router Capable
```

```
Traffic Engineering enabled area
```

```
All capability bits: 0x70000000
```

```
Segment Routing Algorithm TLV: Length: 2
```

```
Algorithm: 0
```

```
Algorithm: 1
```

```
Segment Routing Range TLV: Length: 12
```

```
Range Size: 1000
```

```
SID sub-TLV: Length 3
```

```
Label: 16000
```

# Check OSPF opaque LSAs (2)

## – Check OSPF opaque LSAs

```
RP/0/0/CPU0:R1#show ospf database opaque-area self-originate
```

```
LS age: 454  
Options: (No TOS-capability, DC)  
LS Type: Opaque Area Link  
Link State ID: 7.0.0.1  
Opaque Type: 7  
Opaque ID: 1  
Advertising Router: 172.16.1.1  
LS Seq Number: 80000001  
Checksum: 0x912c  
Length: 44
```

```
Extended Prefix TLV: Length: 20
```

```
Route-type: 1  
AF          : 0  
Flags       : 0x40  
Prefix      : 172.16.1.1/32
```

```
SID sub-TLV: Length: 8
```

```
Flags       : 0x0  
MTID        : 0  
Algo        : 0  
SID Index   : 100
```

# Check CEF entry

## – Check CEF entry

```
RP/0/0/CPU0:R1#show ip cef 172.16.1.6/32
172.16.1.6/32, version 116, internal 0x1000001 0x81 (ptr 0xa1421ff4) [1], 0x0
(0xa13ed4d0), 0xa28 (0xa177207c)
Updated Aug 25 13:29:42.176
local adjacency 10.1.2.2
Prefix Len 32, traffic index 0, precedence n/a, priority 1
via 10.1.2.2/32, GigabitEthernet0/0/0/0, 7 dependencies, weight 0, class 0
[flags 0x0]
  path-idx 0 NHID 0x0 [0xa10fe250 0x0]
  next hop 10.1.2.2/32
  local adjacency
    local label 16600          labels imposed {17600}
  via 10.1.3.3/32, GigabitEthernet0/0/0/1, 7 dependencies, weight 0, class 0
[flags 0x0]
  path-idx 1 NHID 0x0 [0xa10fe2f8 0x0]
  next hop 10.1.3.3/32
  local adjacency
    local label 16600          labels imposed {18600}
```

ECMP

# Check details for R6 prefix SID label

- Check details for R6 prefix SID label

```
RP/0/0/CPU0:R1#show mpls forwarding labels 16006 detai
```

Local Label	Outgoing Label	Prefix or ID	Outgoing Interface	Next Hop	Bytes Switched
-------------	----------------	--------------	--------------------	----------	----------------

16600	17600	SR Pfx (idx 600)	Gi0/0/0/0	10.1.2.2	0
-------	-------	------------------	-----------	----------	---

Updated: Aug 25 13:29:42.177

Version: 116, Priority: 1

Label Stack (Top -> Bottom): { 17600 }

NHID: 0x0, Encap-ID: N/A, Path idx: 0, Backup path idx: 0, Weight: 0

MAC/Encaps: 14/18, MTU: 1500

Packets Switched: 0

18600		SR Pfx (idx 600)	Gi0/0/0/1	10.1.3.3	0
-------	--	------------------	-----------	----------	---

Updated: Aug 25 13:29:42.177

Version: 116, Priority: 1

Label Stack (Top -> Bottom): { 18600 }

NHID: 0x0, Encap-ID: N/A, Path idx: 1, Backup path idx: 0, Weight: 0

MAC/Encaps: 14/18, MTU: 1500

Packets Switched: 0

Traffic-Matrix Packets/Bytes Switched: 0/0

# Ping R6's prefix SID

## – Ping R6's prefix SID

```
RP/0/0/CPU0:R1#RP/0/0/CPU0:R1#ping mpls nil-fec labels 17600 output interface
gigabitEthernet0/0/0/0
```

```
Sending 5, 100-byte MPLS Echos with Nil FEC with labels [17600],
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 rx label,
'P' - no rx intf label prot, 'p' - premature termination of LSP,
'R' - transit router, 'I' - unknown upstream index,
'X' - unknown return code, 'x' - return code 0
```

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

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/6/10 ms
```

# Trace R6's prefix SID

## – Trace R6's prefix SID

```
RP/0/0/CPU0:R1#RP/0/0/CPU0:R1#traceroute mpls nil-fec labels 17600 output
interface gigabitEthernet0/0/0/0
Tracing MPLS Label Switched Path with Nil FEC with labels [17600], timeout is 2
seconds
```

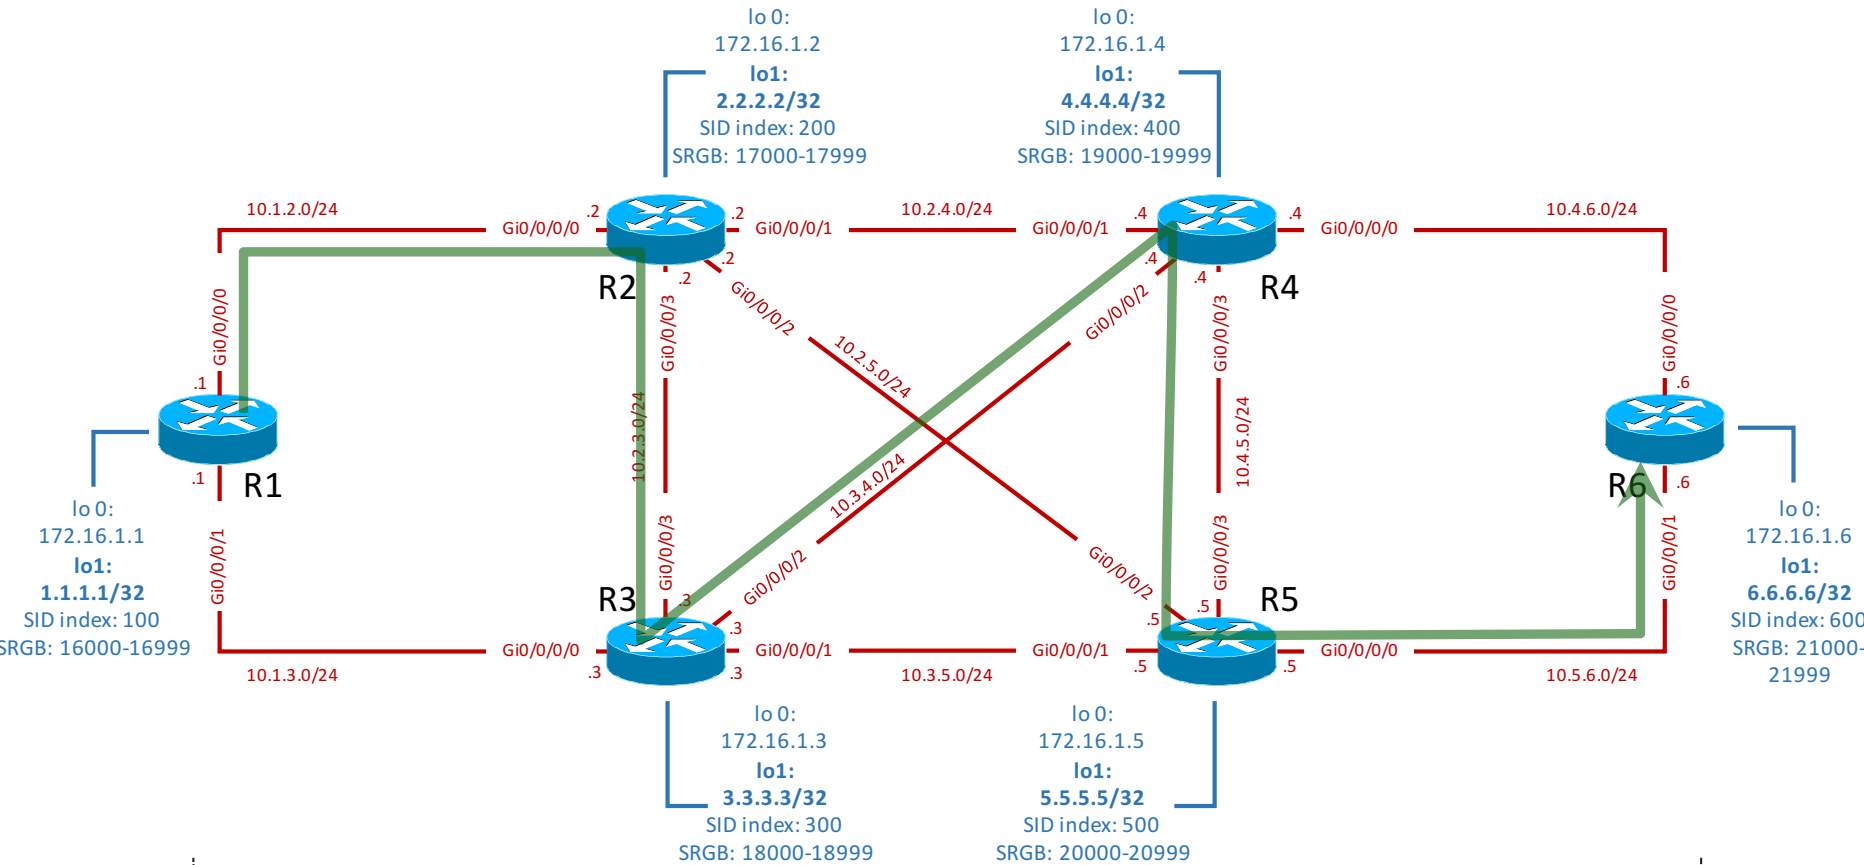
```
Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
'L' - labeled output interface, 'B' - unlabeled output interface,
'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
'M' - malformed request, 'm' - unsupported tlvs, 'N' - no rx label,
'P' - no rx intf label prot, 'p' - premature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

Type escape sequence to abort.

```
0 10.1.2.1 MRU 1500 [Labels: 17600/explicit-null Exp: 0/0]
L 1 10.1.2.2 MRU 1500 [Labels: 20600/explicit-null Exp: 0/0] 0 ms
L 2 10.2.5.5 MRU 1500 [Labels: implicit-null/explicit-null Exp: 0/0] 0 ms
! 3 10.5.6.6 10 ms
```

R2  
R5  
R6

# Part 3: SR-TE



OSPF Instance 1 Area 0.0.0.0



# Create tunnel interface with explicit path

- Create an explicit path.

```
R1(config)# explicit-path name SR-TE-PATH-1
R1(config-expl-path)# index 1 next-address strict ipv4 unicast 172.16.1.2
R1(config-expl-path)# index 2 next-address strict ipv4 unicast 172.16.1.3
R1(config-expl-path)# index 3 next-address strict ipv4 unicast 172.16.1.4
R1(config-expl-path)# index 4 next-address strict ipv4 unicast 172.16.1.5
R1(config-expl-path)# index 5 next-address strict ipv4 unicast 172.16.1.6
```

- Create a TE tunnel interface

```
R1(config)# interface tunnel-te0
R1(config-if)# ipv4 unnumbered Loopback0
R1(config-if)# destination 172.16.1.6
R1(config-if)# path-option 1 explicit name SR-TE-PATH-1 segment-routing
```

# Create loopback interface and static route

- Create a new loopback interface

```
R1(config)# int lo1
R1(config-if)# ipv4 address 1.1.1.1/32
R1(config-if)# exit
```

- Create static route to R6's lo1

```
R1(config)# router static
R1(config-static)# address-family ipv4 unicast
R1(config-static-af)# 6.6.6.6/32 tunnel-te0
R1(config-static-af)# exit
```

# Check tunnel interface

## – Check tunnel interface

```
RP/0/0/CPU0:R1#show mpls traffic-eng tunnels 0 detail

Name: tunnel-te0  Destination: 172.16.1.6  Ifhandle:0xb0
  Signalled-Name: R1_t0
<snip>
  path option 1, (Segment-Routing) type explicit SR-TE-PATH-1 (Basis for
Setup)
  G-PID: 0x0800 (derived from egress interface properties)
<snip>
  Path Info:
  Segment-Routing Path Info (OSPF 1 area 0.0.0.0)
    Segment0[Node]: 172.16.1.2, Label: 16200
    Segment1[Node]: 172.16.1.3, Label: 17300
    Segment2[Node]: 172.16.1.4, Label: 18400
    Segment3[Node]: 172.16.1.5, Label: 19500
    Segment4[Node]: 172.16.1.6, Label: 20600
<snip>
```

# Ping R6's loopback 1 interface

- Ping R6's loopback 1 interface

```
RP/0/0/CPU0: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 = 1/5/9 ms
RP/0/0/CPU0:R1#
```

# Debug of ping to R6

## – Debug of ping to R6

```
RP/0/0/CPU0:Aug 25 14:22:52.138 : netio[309]: mpls_rewrite: #labels 4, out_label
20600, tos 0 eos 1 ttl 255, RA 0, out_intf GigabitEthernet0_0_0_0
RP/0/0/CPU0:Aug 25 14:22:52.138 : netio[309]: mpls_rewrite: out_label[0] 20600,
tos 0 (0x0) eos 1 ttl 255: mpls_header 0x50781ff
RP/0/0/CPU0:Aug 25 14:22:52.138 : netio[309]: mpls_rewrite: out_label[1] 19500,
tos 0 (0x0) eos 0 ttl 255: mpls_header 0x4c2c0ff
RP/0/0/CPU0:Aug 25 14:22:52.138 : netio[309]: mpls_rewrite: out_label[2] 18400,
tos 0 (0x0) eos 0 ttl 255: mpls_header 0x47e00ff
RP/0/0/CPU0:Aug 25 14:22:52.138 : netio[309]: mpls_rewrite: out_label[3] 17300,
tos 0 (0x0) eos 0 ttl 255: mpls_header 0x43940ff
```

# Thank You !

End of session

