Segment Routing Co-existence with LDP

Clarence Filsfils
Kris Michielsen

Segment Routing – Co-existence with LDP

- Co-existence with LDP and other MPLS protocols
- Simple migration from LDP to Segment Routing

Segment Routing and LDP Control Plane Co-existence

Co-existence with other MPLS label distribution protocols

- The MPLS architecture permits concurrent usage of multiple label distribution protocols
  - LDP, RSVP-TE, … and SR control plane can co-exist without interaction
- Each node's Label Manager
  - Reserves a label range (SRGB) for SR control-plane
  - Ensures that all dynamic labels are outside the SRGB block
  - Ensures that a dynamic label is uniquely allocated
- Each LSR must ensure that it can uniquely interpret its incoming labels
  - Adjacency segment: locally unique label allocated by the Label Manager
  - Prefix segment: operator ensures the unique allocation of each label within the allocated SRGB

Segment Routing and LDP Data Plane Co-existence

MPLS-to-MPLS and MPLS-to-IP
label switching and label disposition

- For the MPLS2MPLS and MPLS2IP forwarding entries, SR and LDP can co-exist
  - These entries are indexed on a label
  - The local/incoming labels handled by LDP and SR (or other label distribution protocols) are unique
  - The outgoing label is only significant for the downstream neighbor, not for the local node
  - Multiple MPLS2MPLS and MPLS2IP entries can be programmed for the same prefix
    - > cfr. LSP midpoint cross-connect

All nodes: SR + LDP SR LDP

MPLS-to-MPLS and MPLS-to-IP

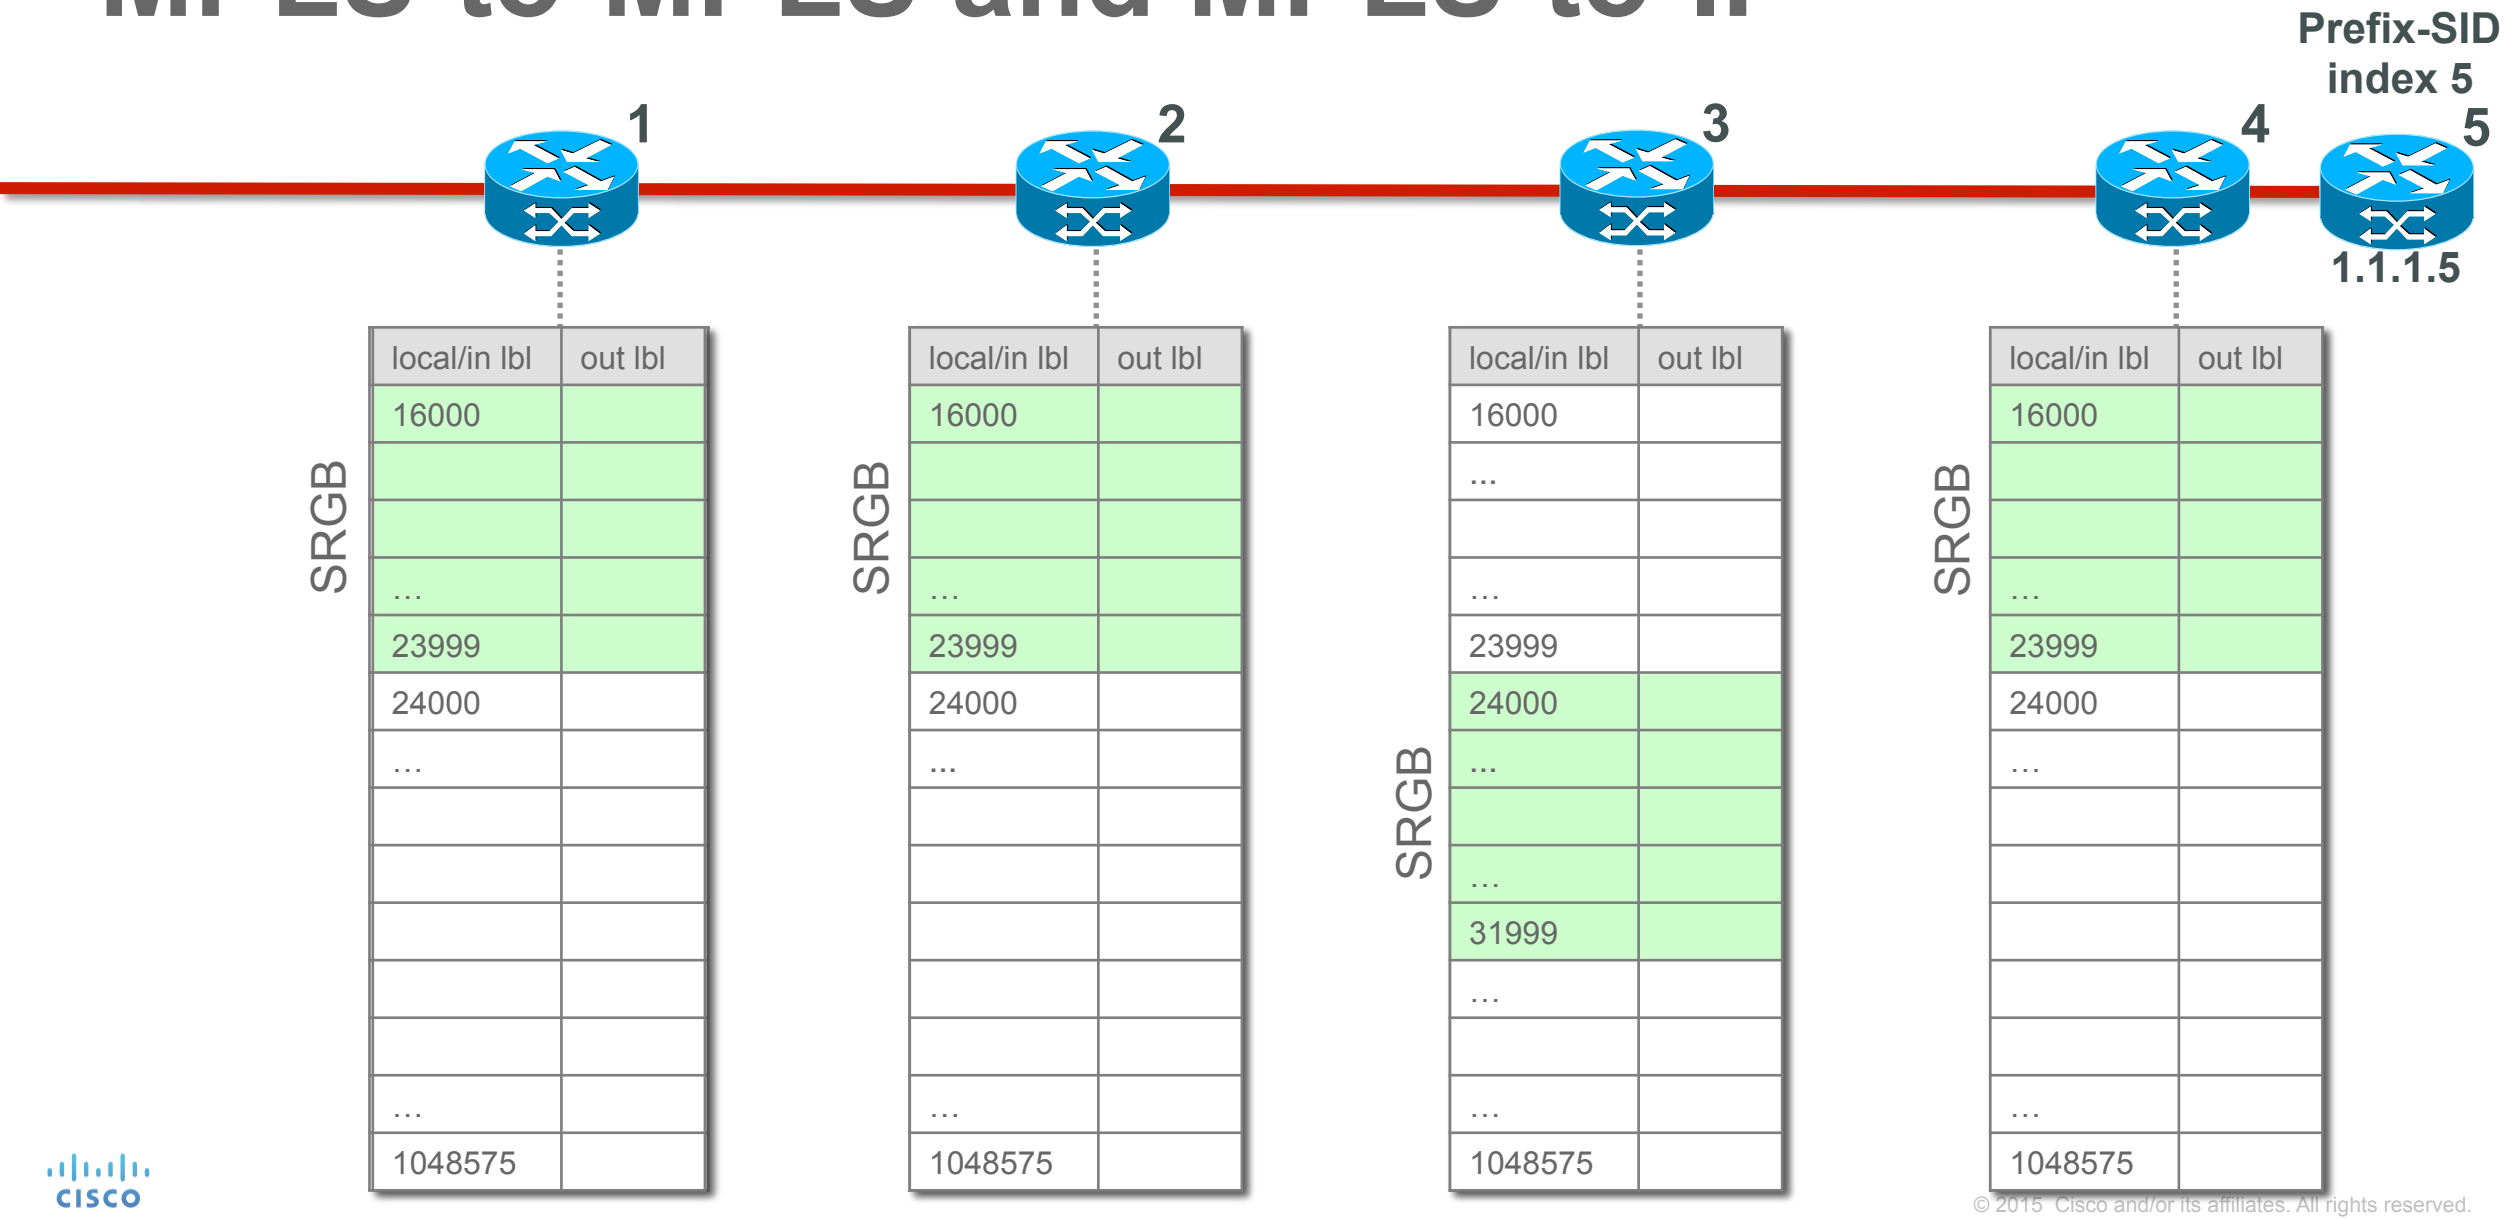

All nodes: SR + LDP SR LDP

MPLS-to-MPLS and MPLS-to-IP

SR Prefix Segment to 1.1.1.5/32

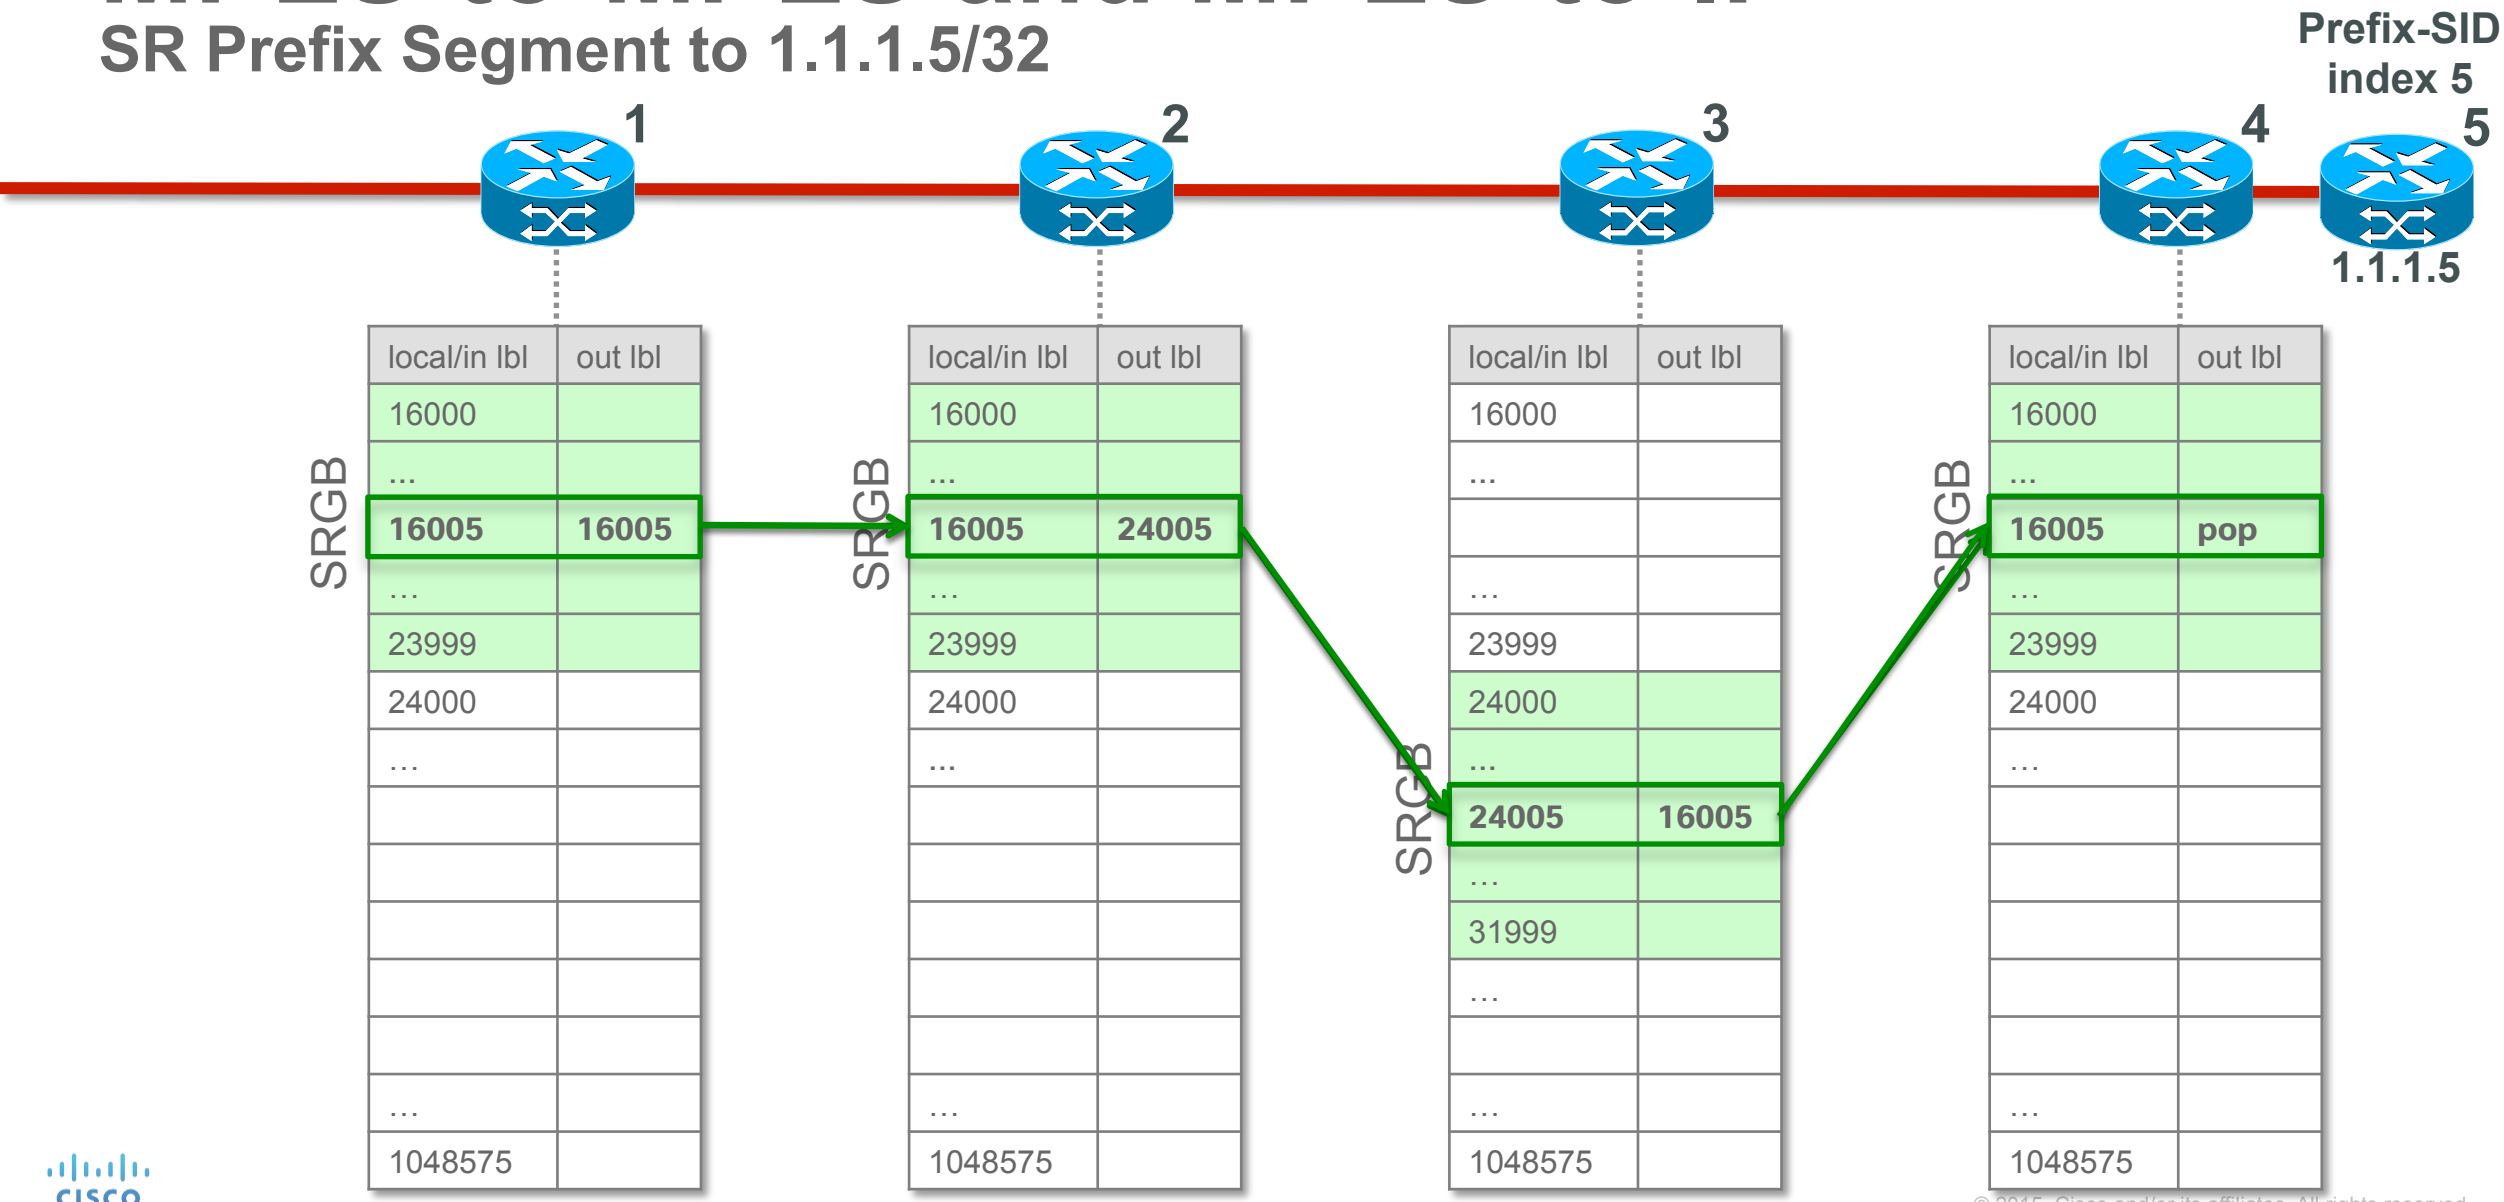

All nodes: SR + LDP

SR LDP

MPLS-to-MPLS and MPLS-to-IP

LDP FEC to 1.1.1.5/32

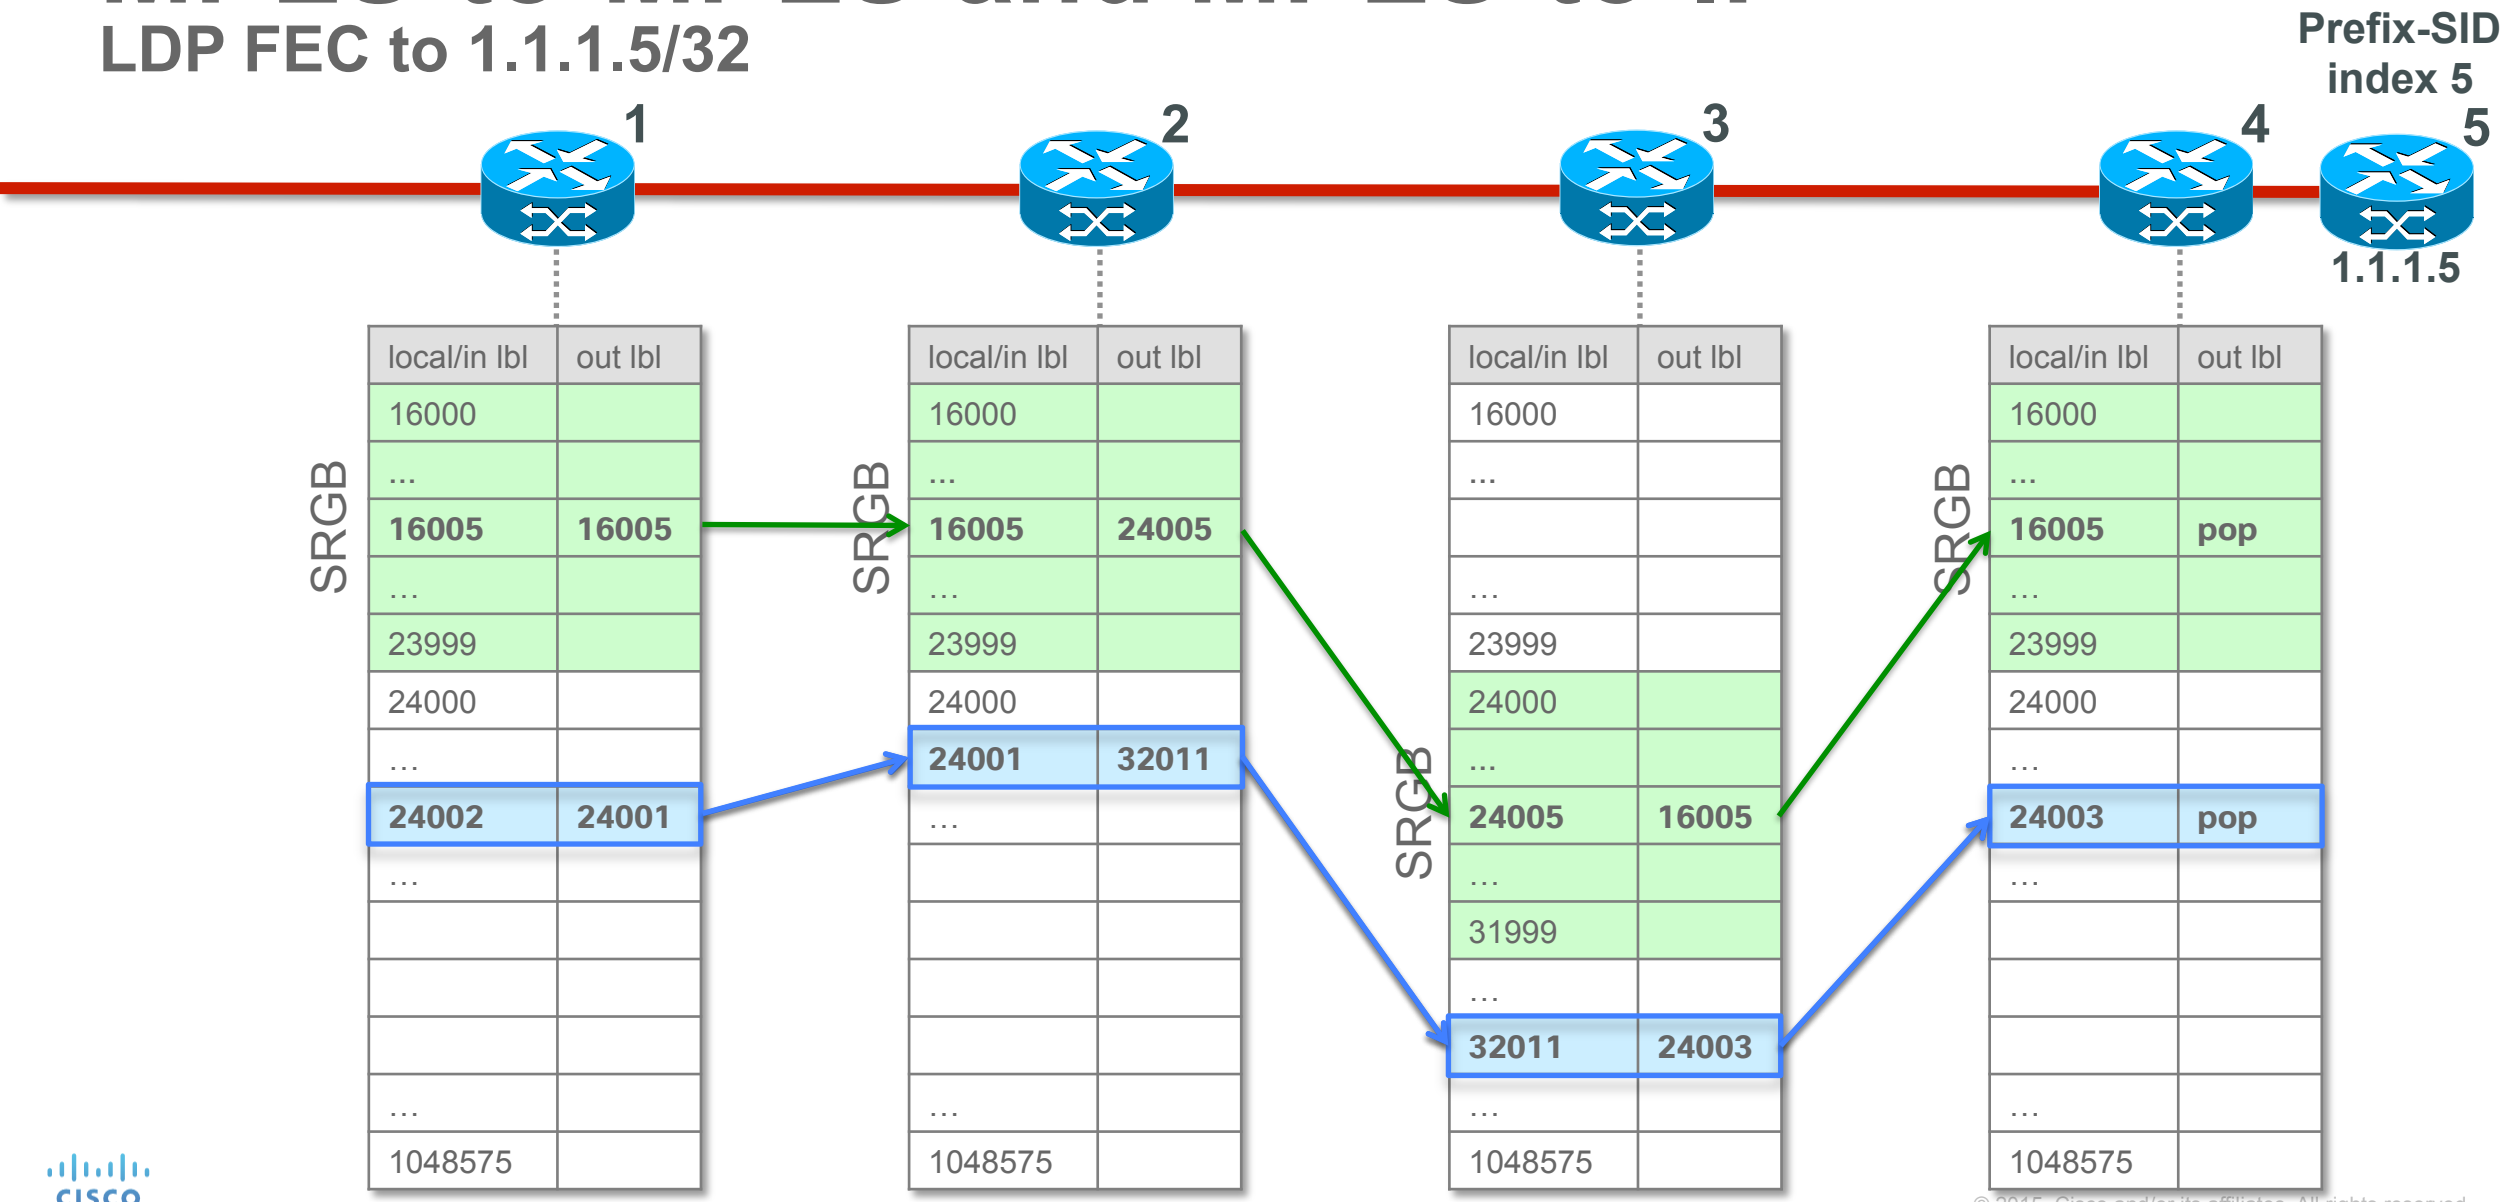

IP-to-MPLS – label imposition

- Multiple IP2MPLS entries (e.g. LDP and SR) for the same prefix path cannot co-exist
  - These label imposition forwarding entries are indexed on the prefix
  - A forwarding table lookup returns one or more paths to the destination
  - Each path has a single IP2MPLS entry programmed
  - If multiple paths lead to the destination, each path has its own IP2MPLS entry
    - E.g. one path imposing an LDP label, another path imposing an SR label

All nodes: SR + LDP

SR LDP

IP-to-MPLS: which label must be imposed?

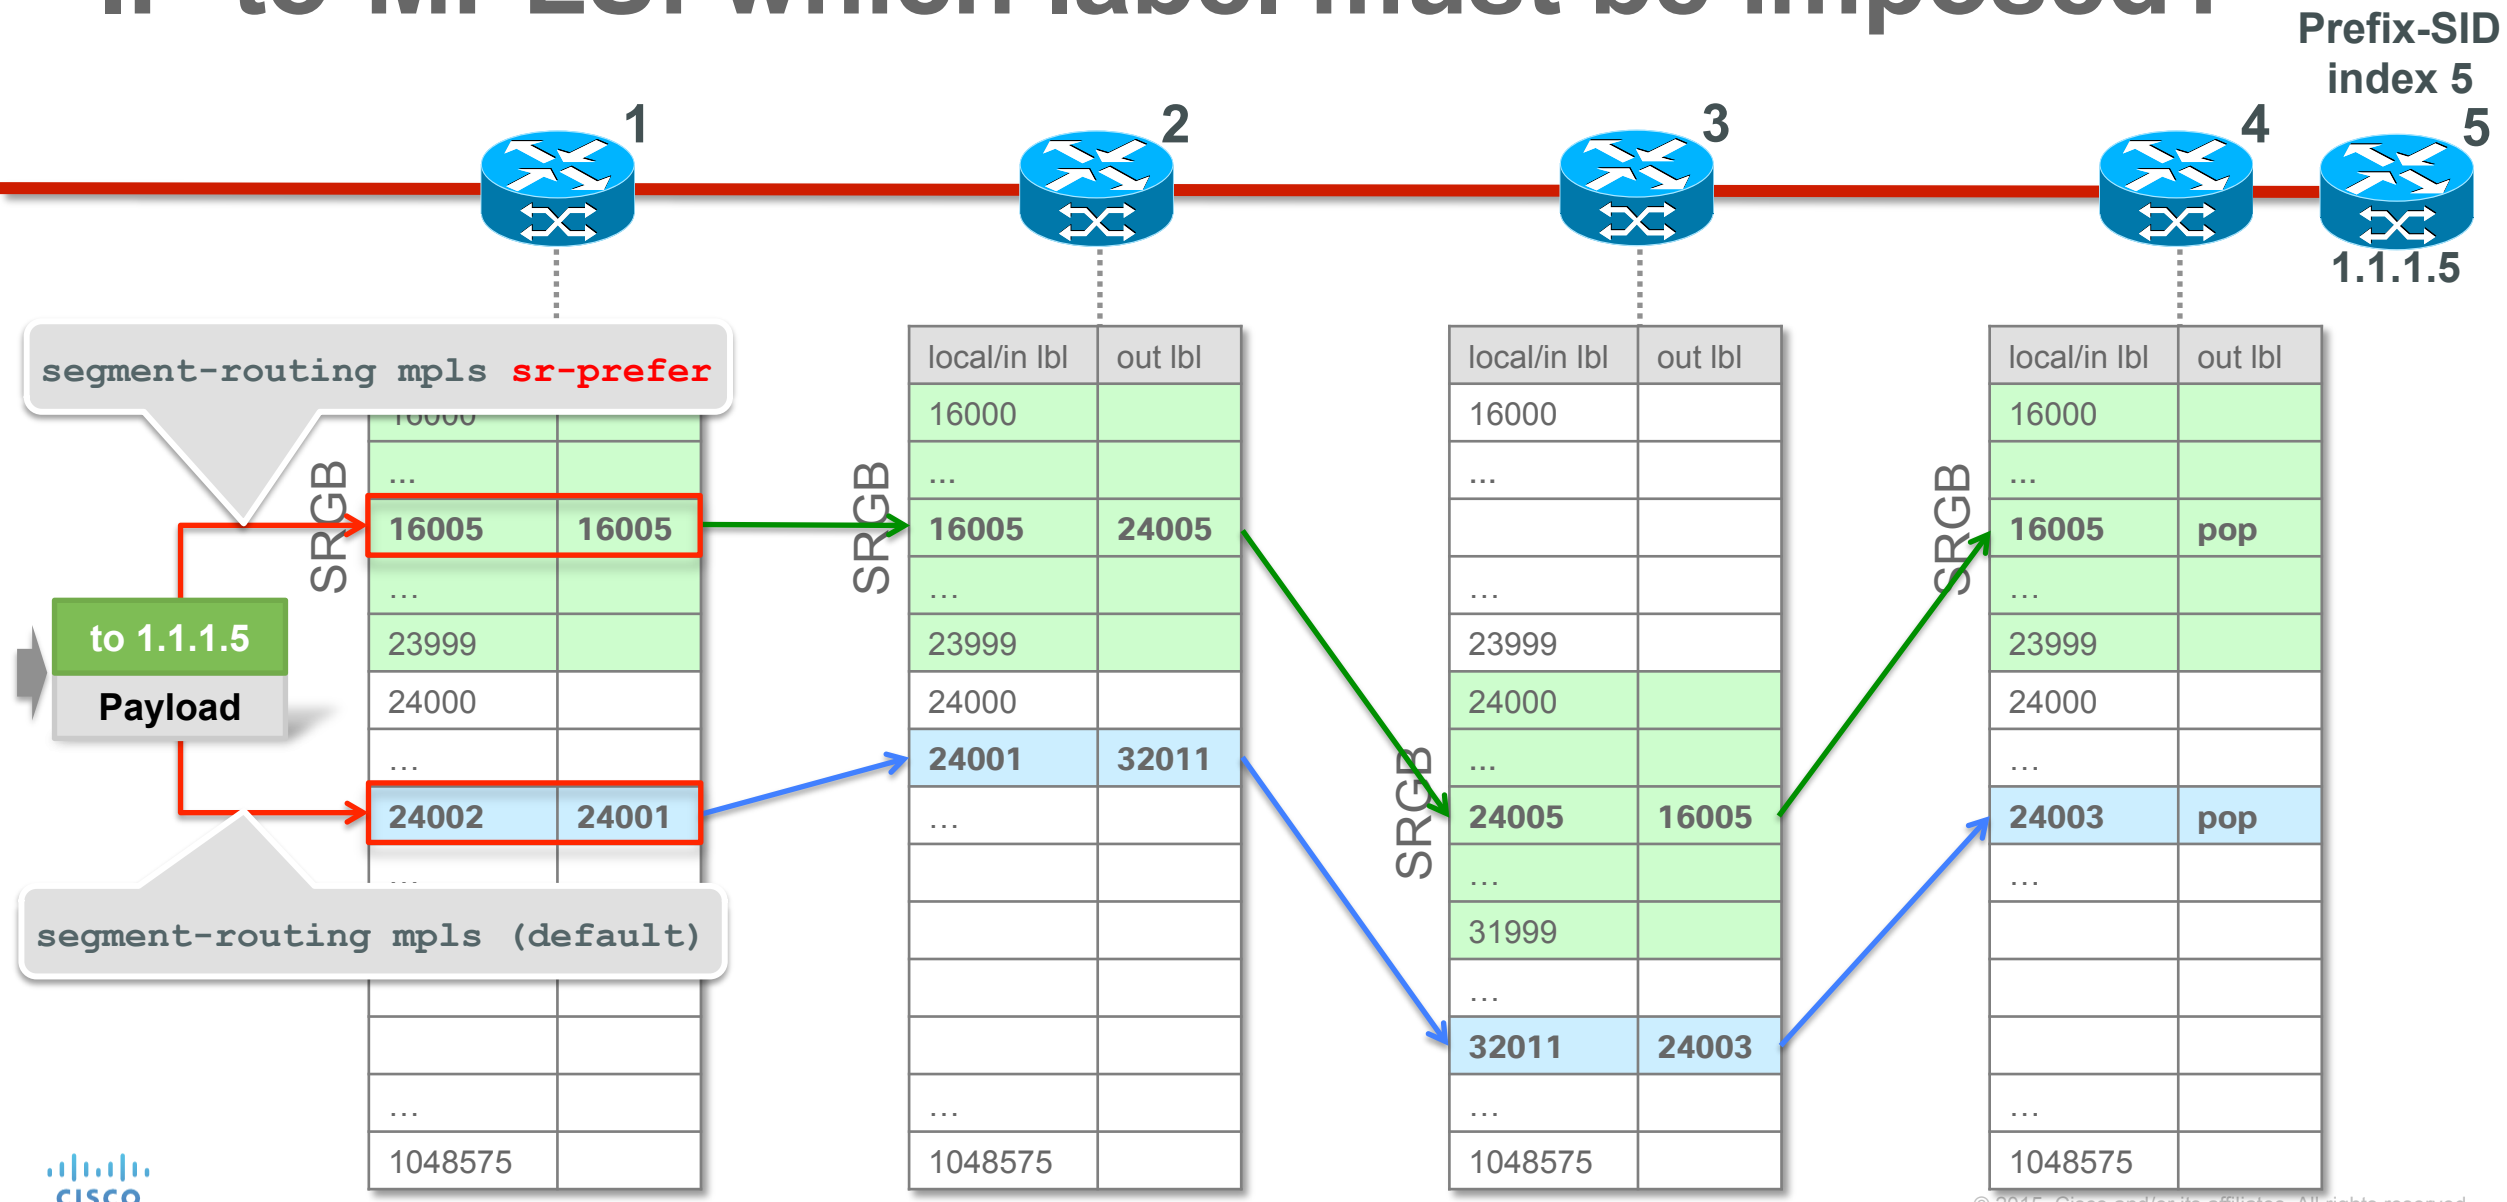

IP-to-MPLS – label imposition

- For IP2MPLS forwarding, LDP XOR SR entry can be inserted into FIB
  - Only one IP2MPLS entry can exists for each prefix path
- Default: LDP label imposition is preferred

```
router isis 1
 address-family ipv4|6 unicast
  segment-routing mpls sr-prefer
```

```
router ospf 1
 segment-routing mpls
 segment-routing sr-prefer
```

IGP/SR and LDP programming FIB

- This diagram illustrates the behavior of node1 on slide 8

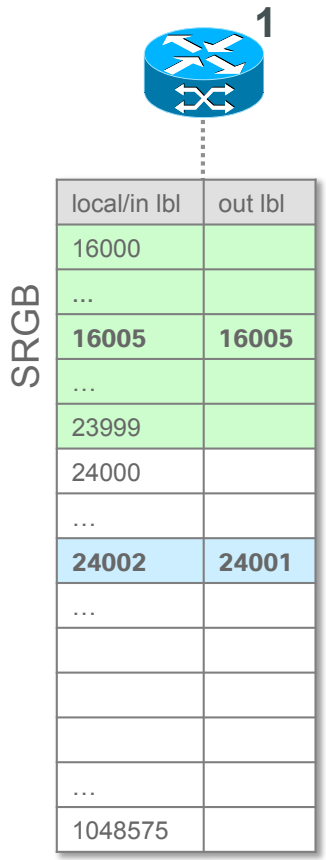

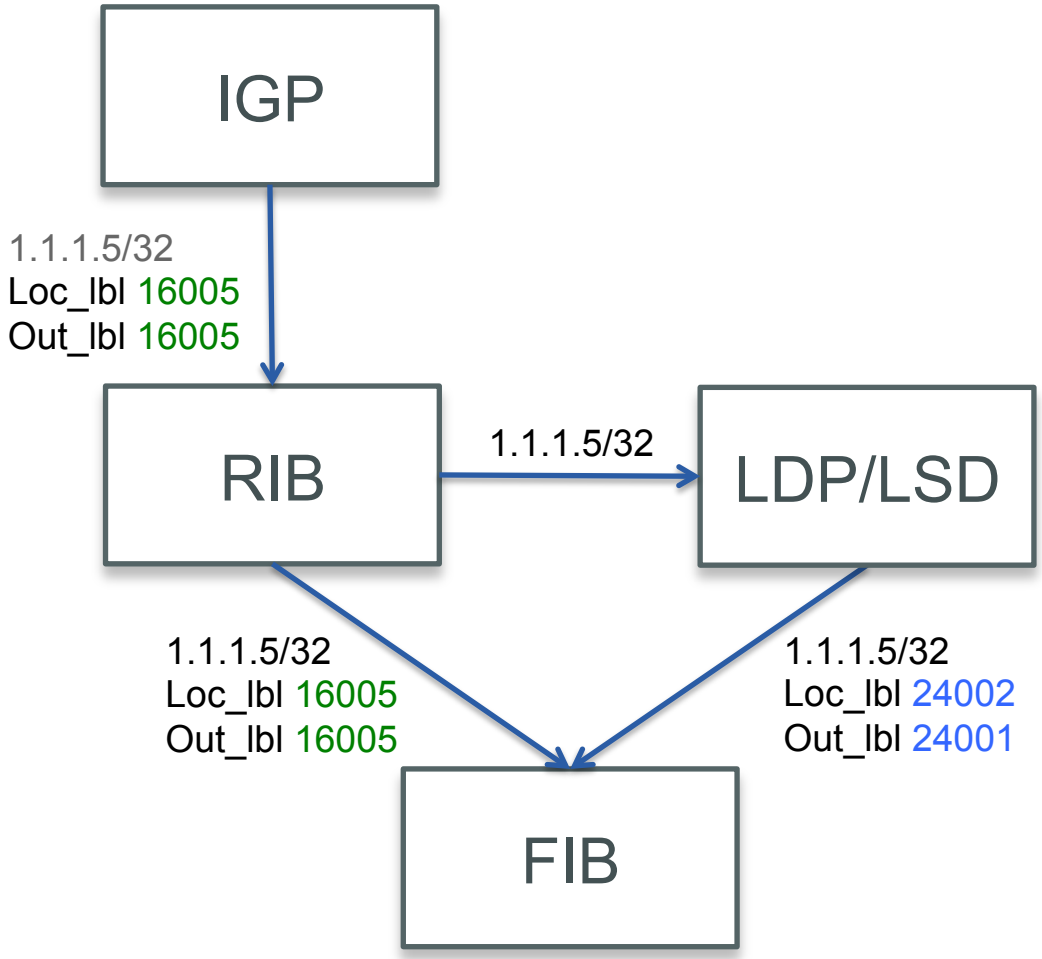

Loc_lbl: local label, allocated by local node
Out_lbl: outgoing label

16005: SR label
24002: LDP label

IGP/SR and LDP programming FIB

- This diagram illustrates the default behavior: prefer LDP label imposition

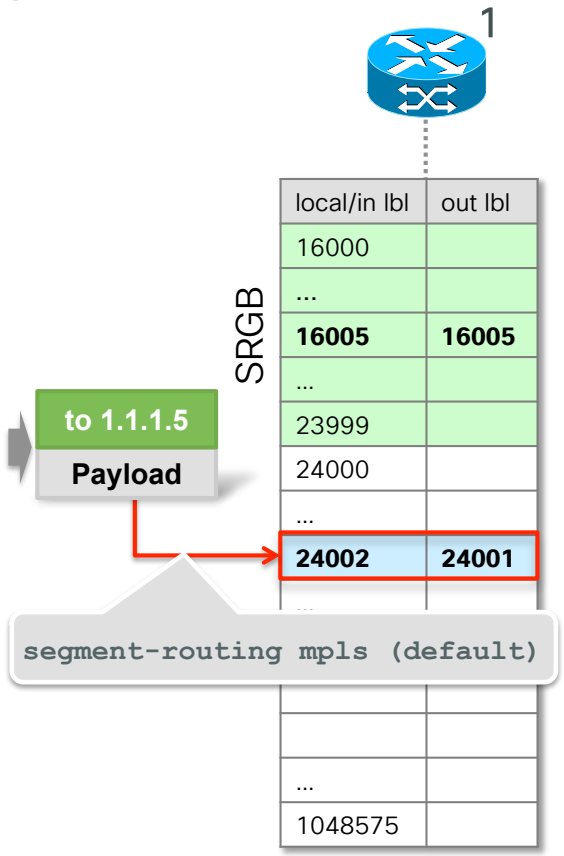

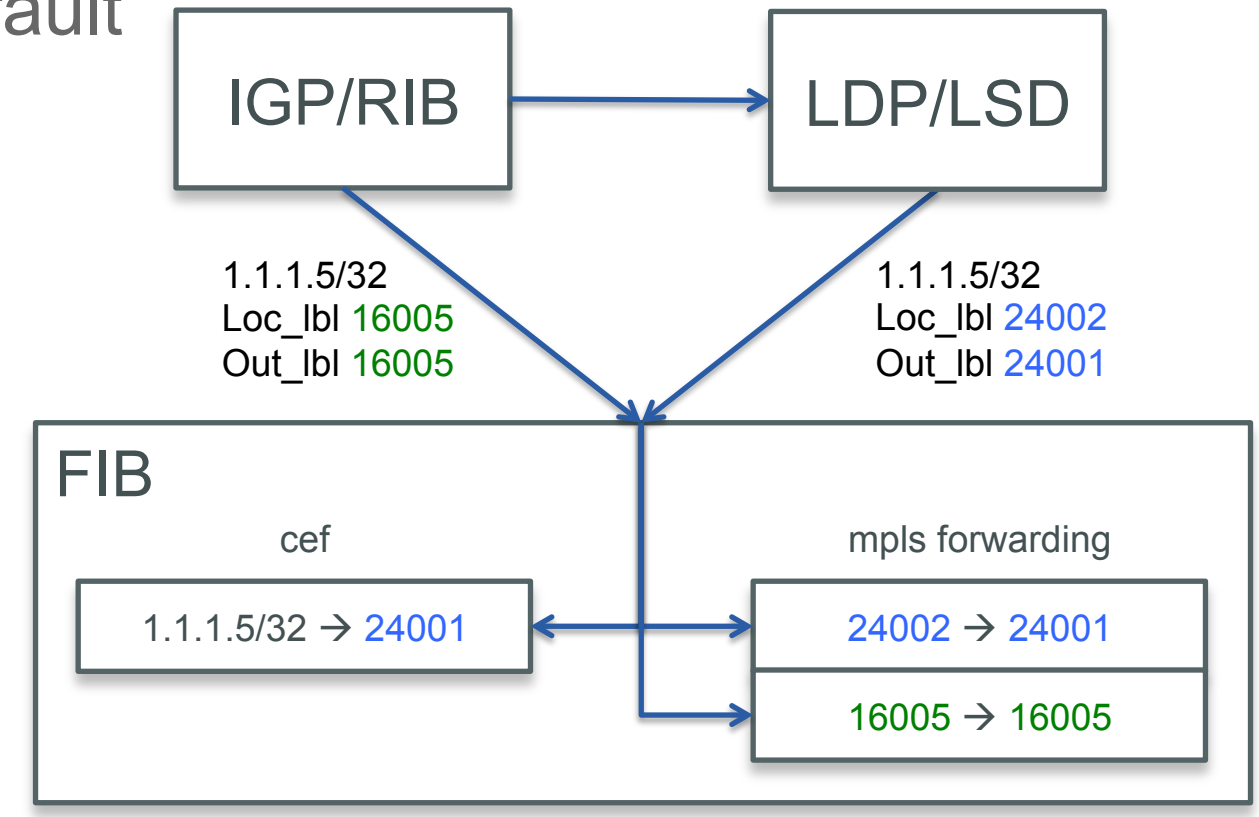

Loc_lbl: local label, allocated by local node
Out_lbl: outgoing label

16005: SR label
24002: LDP label

IGP/SR and LDP programming FIB

- This diagram illustrates the behavior when preferring SR label imposition

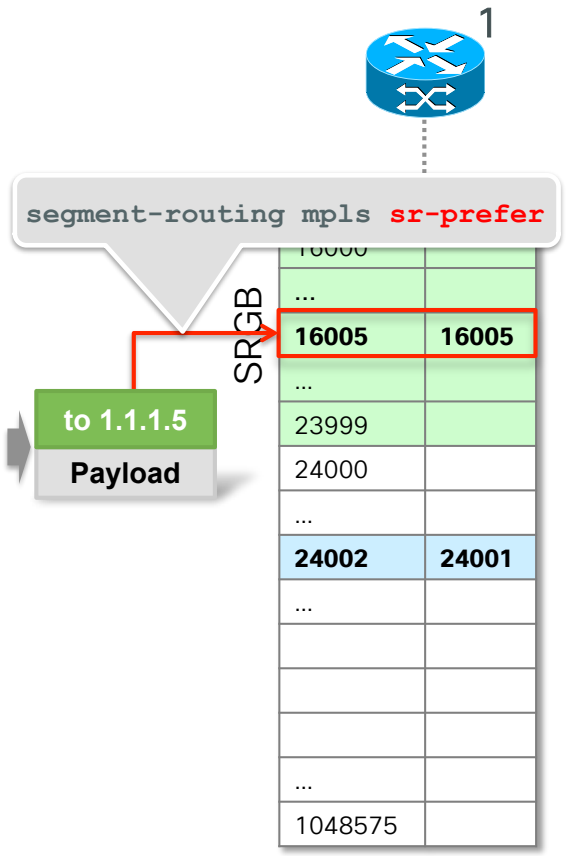

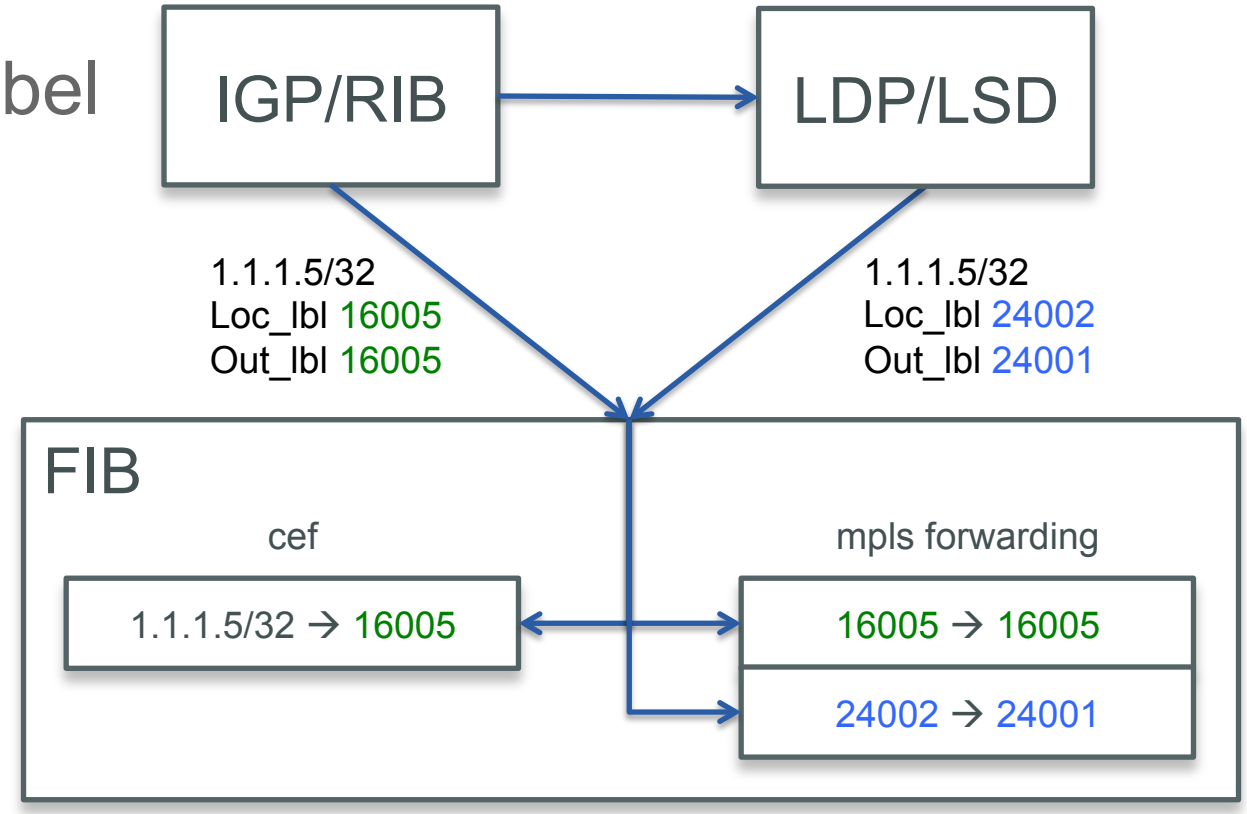

Loc_lbl: local label, allocated by local node
Out_lbl: outgoing label

16005: SR label
24002: LDP label

All nodes: SR + LDP

MPLS-to-MPLS forwarding entries


```
RP/0/0/CPU0:xrvr-3#show mpls forwarding labels 24003
Local  Outgoing    Prefix            Outgoing     Next Hop        Bytes
Label  Label       or ID             Interface                    Switched
------ ----------- ----------------- ------------ --------------- ------------
24003  24005       100.0.0.3/32      Gi0/0/0/0    99.2.3.3        5670


RP/0/0/CPU0:xrvr-3#show mpls forwarding labels 16001
Local  Outgoing    Prefix            Outgoing     Next Hop        Bytes
Label  Label       or ID             Interface                    Switched
------ ----------- ----------------- ------------ --------------- ------------
16001  16001       SR Pfx (idx 1)    Gi0/0/0/0    99.2.3.3        345243
```

MPLS-to-MPLS **LDP** label entry

Both entries are present regardless of the preference setting

MPLS-to-MPLS **SR** label entry

IP-to-MPLS forwarding entries with LDP preference

All nodes: SR + LDP

default (LDP is preferred)

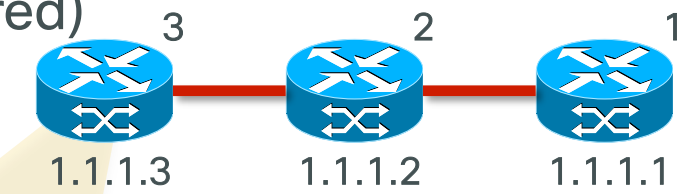


```
RP/0/0/CPU0:xrvr-3#show mpls ldp bindings 1.1.1.1/32 neighbor 1.1.1.2
1.1.1.1/32, rev 24
        Local binding: label: 24003
        Remote bindings: (1 peer)
            Peer                 Label
            -----------------    ---------
            1.1.1.2:0            24005

RP/0/0/CPU0:xrvr-3#show cef 1.1.1.1/32
1.1.1.1/32, version 222, internal 0x4000001 0x0 (ptr 0xa1376074) [1], 0x0 (0xa135b560),
0x228 (0xa1411118)
 Updated May 21 07:08:50.475
 local adjacency 99.2.3.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via 99.2.3.2, GigabitEthernet0/0/0/0, 9 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa0e300bc 0x0]
    next hop 99.2.3.2
    tx adjacency
     local label 24003      labels imposed {24005}
```

Local LDP label

Neighbor's (remote) LDP label

IP-to-MPLS cef entry has LDP labels programmed

SR and LDP – IP-to-MPLS with SR preference

All nodes: SR + LDP

sr-prefer configured


```
RP/0/0/CPU0:xrvr-3#show route 1.1.1.1/32 detail
Routing entry for 1.1.1.1/32
  Known via "isis 1", distance 115, metric 2, type level-2
  Installed May 21 07:08:45.345 for 00:35:05
  Routing Descriptor Blocks
    99.2.3.2, from 1.1.1.1, via GigabitEthernet0/0/0/0
      Route metric is 2
      Label: 0x3e81 (16001)
      Tunnel ID: None
      Extended communities count: 0
      Path id:1       Path ref count:0
      NHID:0x1(Ref:6)
  Route version is 0xa (10)
  Local Label: 0x3e81 (16001)
<...>

RP/0/0/CPU0:xrvr-3#show cef 1.1.1.1/32
1.1.1.1/32, version 222, internal 0x4000001 0x0 (ptr 0xa1376074) [1], 0x0 (0xa135b560),
0x228 (0xa1411118)
 Updated May 21 07:08:50.475
 local adjacency 99.2.3.2
 Prefix Len 32, traffic index 0, precedence n/a, priority 3
   via 99.2.3.2, GigabitEthernet0/0/0/0, 9 dependencies, weight 0, class 0 [flags 0x0]
    path-idx 0 NHID 0x0 [0xa0e300bc 0x0]
    next hop 99.2.3.2
    tx adjacency
     local label 16001      labels imposed {16001}
```

Neighbor's (remote) SR label

Local SR label

IP-to-MPLS cef entry has SR labels programmed

Segment Routing and LDP “Ships in the night” Deployment Model

“Ships in the Night” Deployment Model

- LDP and SR are kept independent
  - continuous SR connectivity between SR PEs required
  - continuous LDP connectivity between LDP PEs required
  - no SR to LDP or LDP to SR interworking required
- Other deployment models are possible: see “SR/LDP interworking” section

Simplest migration LDP to SR

Assumptions:
- all the nodes can be upgraded to SR
- all the services can be upgraded to SR

- Initial state: All nodes run LDP, not SR

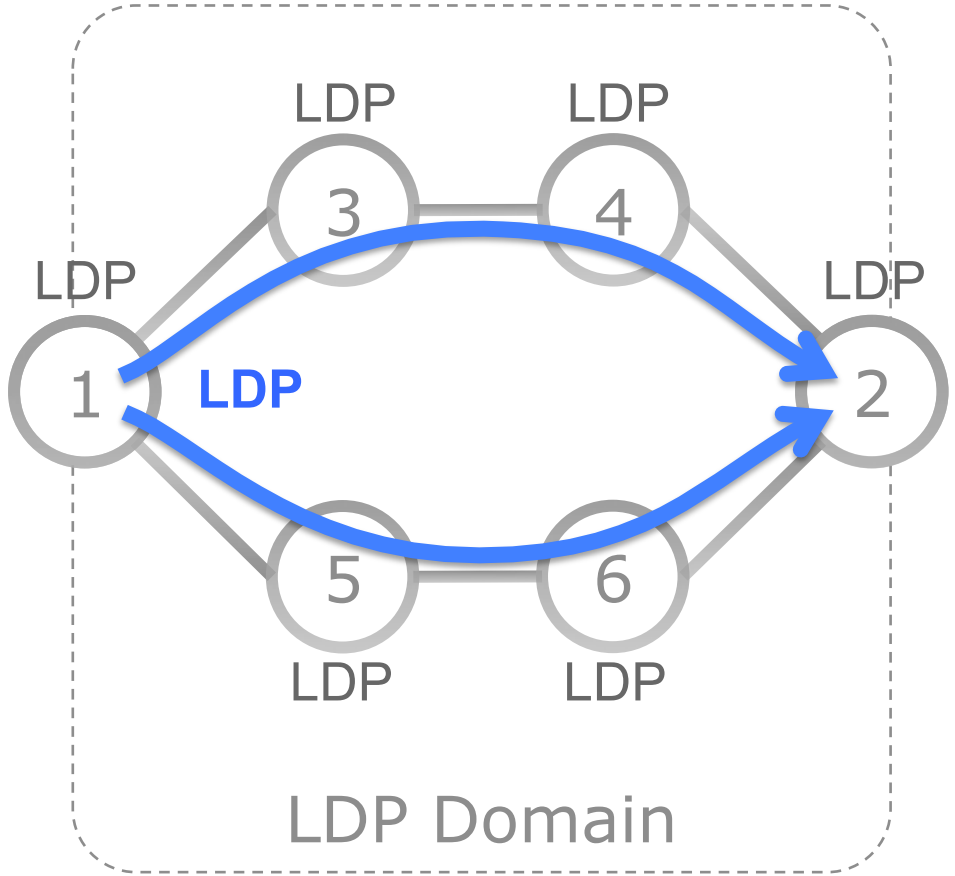

Simplest migration LDP to SR

Assumptions:
- all the nodes can be upgraded to SR
- all the services can be upgraded to SR

- Initial state: All nodes run LDP, not SR
- Step1: All nodes are upgraded to SR
  - In no particular order
  - leave default LDP label imposition preference

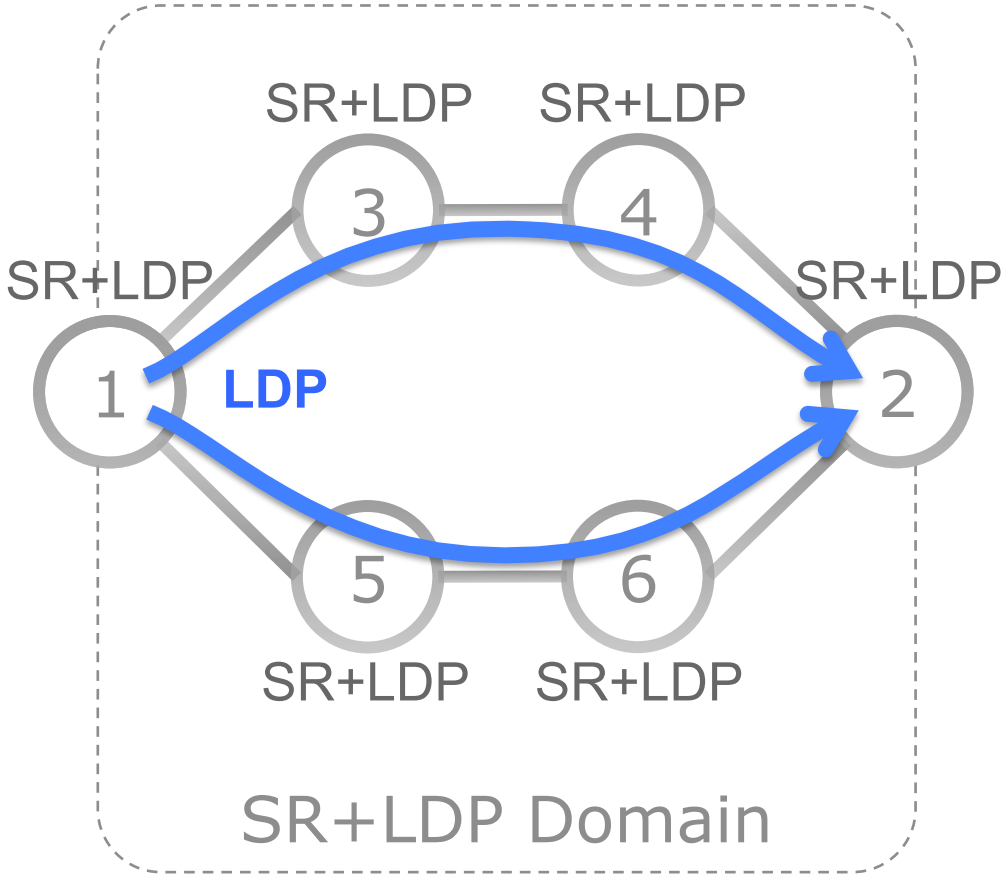

Simplest migration LDP to SR

Assumptions:
- all the nodes can be upgraded to SR
- all the services can be upgraded to SR

- Initial state: All nodes run LDP, not SR
- Step1: All nodes are upgraded to SR
  - In no particular order
  - leave default LDP label imposition preference
- Step2: All PEs are configured to prefer SR label imposition
  - In no particular order

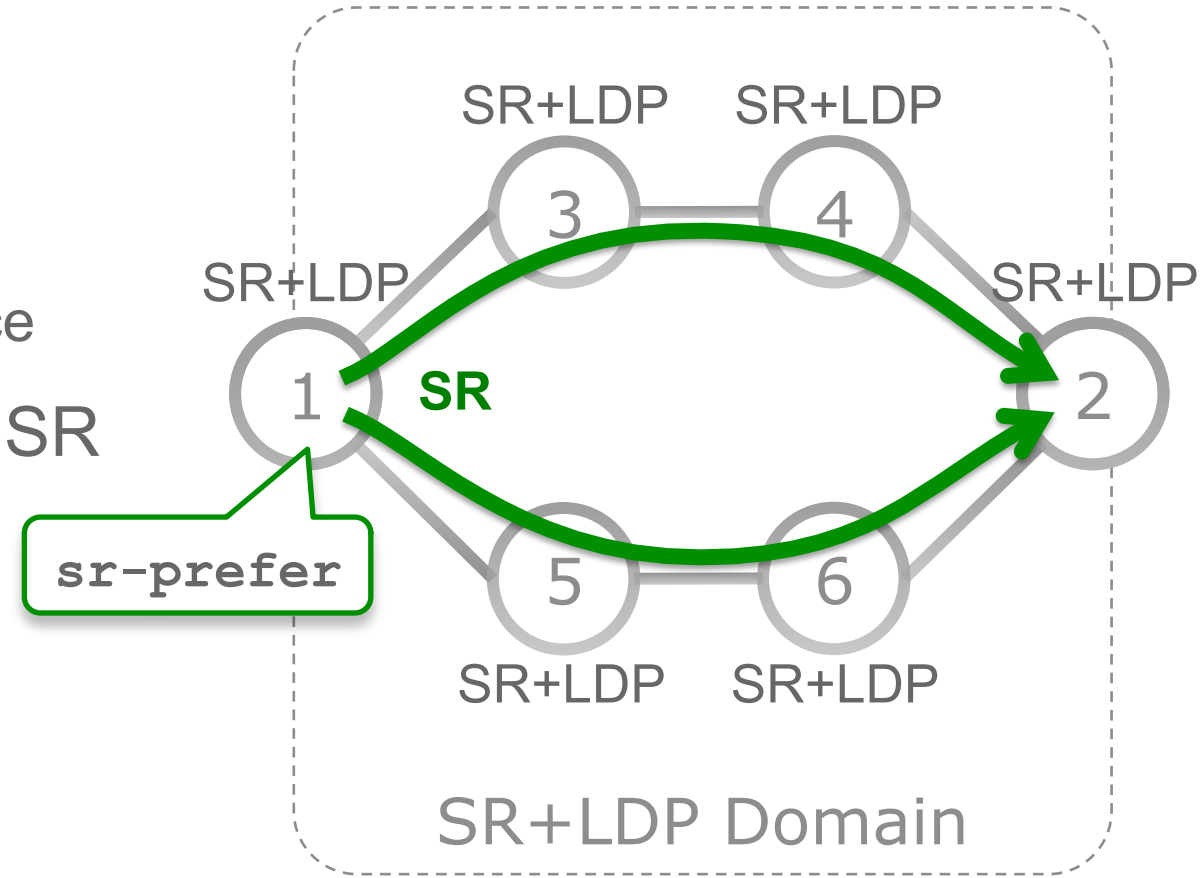

Simplest migration LDP to SR

Assumptions:
- all the nodes can be upgraded to SR
- all the services can be upgraded to SR

- Initial state: All nodes run LDP, not SR
- Step1: All nodes are upgraded to SR
  - In no particular order
  - leave default LDP label imposition preference
- Step2: All PEs are configured to prefer SR label imposition
  - In no particular order
- Step3: LDP is removed from the nodes in the network
  - In no particular order
- Final state: All nodes run SR, not LDP

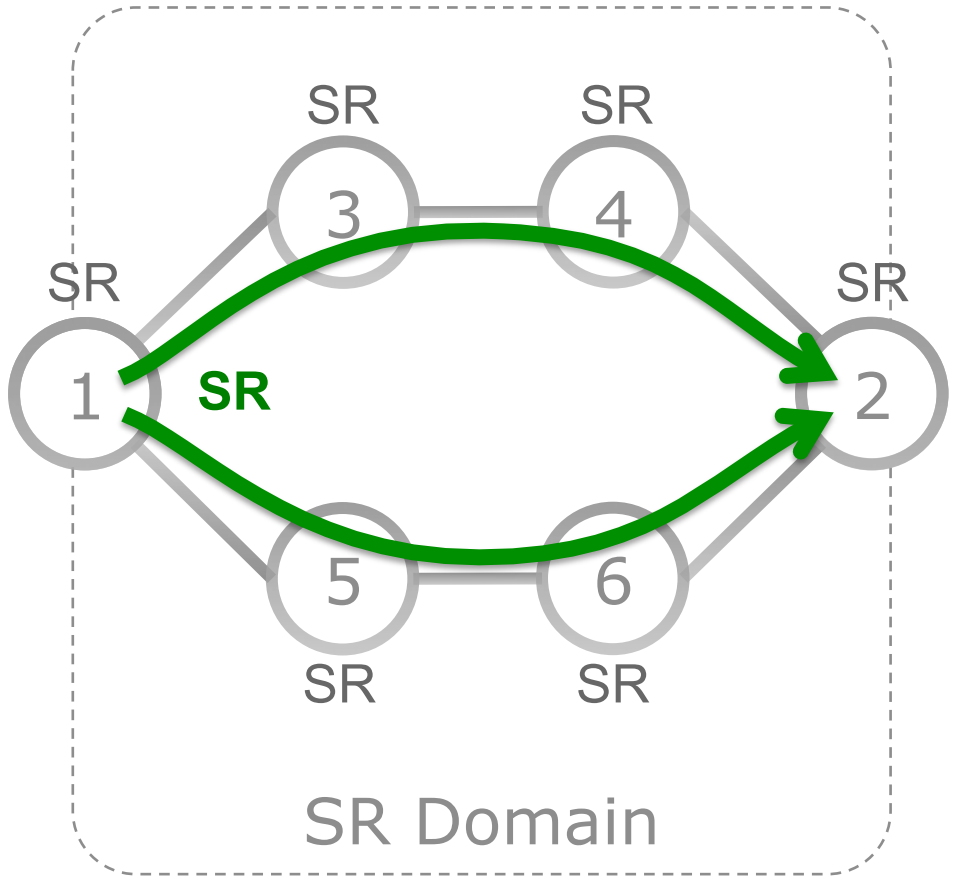

Visit us:

cisco.com

segment-routing.net

Acknowledgements:

Ahmed Bashandy
Robert Hanzl
Steven Luong
Stefano Previdi
Peter Psenak

Thank you.