Segment Routing IGP Control Plane

Clarence Filsfils
Kris Michielsen

Segment Routing – IGP Control plane

- Using IS-IS or OSPF to distribute segments
- Configuring Segment Routing under IGP
- Segment Routing in a multi-area, multi-level network
- Verifying Segment Routing advertisements

Segment Routing IS-IS Overview

SR IS-IS Control Plane Overview

- IS-IS Segment Routing functionality
  - IPv4 and IPv6 control plane
  - Level 1, level 2 and multi-level routing
  - Prefix Segment ID (Prefix-SID) for host prefixes on loopback interfaces
  - Adjacency Segment IDs (Adj-SIDs) for adjacencies
    - >Non-protected adj-SIDs and protected (since IOS XR 5.3.2) adj-SIDs
      - See SR-TE presentation for more information
  - Prefix-to-SID mapping advertisements (mapping server)
  - MPLS penultimate hop popping (PHP) and explicit-null signaling

IS-IS TLV Extensions

- SR for IS-IS introduces support for the following (sub-)TLVs:

– SR Capability sub-TLV (2)	IS-IS Router Capability TLV (242)
– Prefix-SID sub-TLV (3)	Extended IP reachability TLV (135)
– Prefix-SID sub-TLV (3)	IPv6 IP reachability TLV (236)
– Prefix-SID sub-TLV (3)	Multitopology IPv6 IP reachability TLV (237)
– Prefix-SID sub-TLV (3)	SID/Label Binding TLV (149)
– Adjacency-SID sub-TLV (31)	Extended IS Reachability TLV (22)
– LAN-Adjacency-SID sub-TLV (32)	Extended IS Reachability TLV (22)
– Adjacency-SID sub-TLV (31)	Multitopology IS Reachability TLV (222)
– LAN-Adjacency-SID sub-TLV (32)	Multitopology IS Reachability TLV (222)
– SID/Label Binding TLV (149)	

- Implementation based on *draft-ietf-isis-segment-routing-extensions-02*

Segment Routing OSPF Overview

SR OSPF Control Plane Overview

- OSPF Segment Routing functionality
  - OSPFv2 control plane
  - Multi-area
  - IPv4 Prefix Segment ID (Prefix-SID) for host prefixes on loopback interfaces
  - Adjacency Segment ID (Adj-SIDs) for adjacencies
    - >Non-protected adj-SIDs and protected (since OSPF SR-TE release) adj-SIDs
  - MPLS penultimate hop popping (PHP) and explicit-null signaling

OSPF Extensions

- OSPF adds to the Router Information Opaque LSA (type 4):
  - SR-Algorithm TLV (8)
  - SID/Label Range TLV (9)
- OSPF defines new Opaque LSAs to advertise the SIDs
  - OSPFv2 Extended Prefix Opaque LSA (type 7)
    - > OSPFv2 Extended Prefix TLV (1)
      - Prefix SID Sub-TLV (2)
  - OSPFv2 Extended Link Opaque LSA (type 8)
    - > OSPFv2 Extended Link TLV (1)
      - Adj-SID Sub-TLV (2)
      - LAN Adj-SID Sub-TLV (3)
- Implementation is based on

  draft-ietf-ospf-prefix-link-attr-01 and draft-ietf-ospf-segment-routing-extensions-02

Configuring Segment Routing

IS-IS Segment Routing Configuration

```
router isis 1
 address-family ipv4|ipv6 unicast
  metric-style wide
  segment-routing|mpls
 !
```

enable SR capability

enable SR for the MPLS data plane

- MPLS forwarding is enabled on all non-passive IS-IS interfaces
- Adjacency-SIDs are allocated and distributed for all adjacencies
  - Non-protected adj-SIDs and protected (since IOS XR 5.3.2) adj-SIDs
  - See SR-TE section

IS-IS Segment Routing Configuration

```
router isis 1
 address-family ipv6 unicast
  metric-style wide
  segment-routing ipv6
 !
```

enable SR capability

enable SR for the IPv6 extension-header data plane

- SRv6 Extension Header data plane is outside the scope of this presentation

OSPF Segment Routing Configuration
Recommended

```
router ospf 1
 segment-routing mpls
 segment-routing forwarding mpls
```

In a later release, SR forwarding will be enabled by default. This config line will no longer be required. (CSCuw93707)

- **`segment-routing forwarding mpls`** must be configured to install SIDs – received by OSPF – in the forwarding table
- MPLS forwarding is enabled on all **`segment-routing forwarding`** enabled OSPF interfaces
- Adjacency-SIDs are allocated and distributed for **`segment-routing forwarding`** enabled adjacencies
- Configuration under ospf instance is recommended, but can be customized
  - See next slides

OSPF configuration inheritance – Reminder

- OSPF configuration has a hierarchical model:
  - Instance
    - > Area
      - Interface

```
router ospf 1
 area 0
  interface GigabitEthernet0/0/0/0
```

- OSPF configurations on a higher level in the hierarchy are inherited by the lower level(s)
  - Ex1: OSPF interface configuration (e.g. **`hello-interval`**) configured at the instance level is inherited by all interfaces of the OSPF instance
  - Ex2: OSPF interface configuration (e.g. **`hello-interval`**) configured at the area level is inherited by all interfaces in the area

OSPF configuration inheritance – Reminder

- OSPF Configuration on a lower level in the hierarchy (more specific) is preferred over a higher level configuration
  - Ex: OSPF configuration on an interface is preferred over the configuration done at the instance or area level

```
router ospf 1
 hello-interval 10        !! OSPF interfaces use hello-interval 10
 area 0
  hello-interval 20       !! Area 0 interfaces use hello-interval 20
  interface GigabitEthernet0/0/0/0
   hello-interval 30      !! This interface uses hello-interval 30
```

OSPF configuration inheritance – Reminder

- OSPF Configurations usually have a **`disable`** keyword
  - Ex: Enable an OSPF interface configuration on instance level, then disable it on interface level for specific interfaces

```
router ospf 1
 bfd fast-detect                 !! Enable on instance level
 area 0
  bfd fast-detect disable        !! Disable on area level
 !
 area 1
  interface GigabitEthernet0/0/0/0
   bfd fast-detect disable       !! Disable on interface level
```

OSPF Segment Routing Configuration

```
router ospf 1
 area 0
  segment-routing mpls                    !! Area command
  segment-routing forwarding mpls         !! Interface command
  interface GigabitEthernet0/0/0/0
   segment-routing forwarding disable !! Interface command
```

- **`segment-routing mpls`** is an ospf area command, can be applied per area
  - Ospf inheritence rules are applicable
- **`segment-routing forwarding mpls`** is an ospf interface command, can be applied per interface
  - Ospf inheritence rules are applicable
- In the example, SR is enabled for all interfaces in area0, except Gi0/0/0/0

Segment Routing Global Block (SRGB)

Segment Routing Global Block

- Default SRGB is [16,000-23,999]
  - Default SRGB configuration not shown in configuration
- Non-default SRGB can be configured per IGP instance
- Multiple IGP instances can use the same SRGB or use different non-overlapping SRGBs
- Segment Routing Global Block can be configured in global configuration (IOS XR 6.0)
  - SRGB under IGP instance has precedence over SRGB in global configuration
- See “Segment Routing Global Block” slides for more details

Segment Routing Global Block (SRGB)
Example

```
segment-routing
 global-block 18000 19999
!
router ospf 1
 segment-routing mpls
 !! no segment-routing global-block config
```

Configure a non-default global SRGB 18,000 – 19,999

```
RP/0/0/CPU0:xrvr-1#show mpls label table detail
Table Label   Owner                              State
----- ------- ---------------------------------- ------
<...snip...>
0     18000   OSPF(A):ospf-1                     InUse  No
  (Lbl-blk SRGB, vers:0, (start_label=18000, size=2000)
<...snip...>
```

OSPF SRGB

Start_label = 18,000

Size = 2,000

Non-default SRGB label block allocation for OSPF
[18,000 – 19,999]

Segment Routing Global Block (SRGB)
Example

```
!! no global segment-routing global-block config
router isis 1
 segment-routing mpls
 segment-routing global-block 18000 19999
```

Configure an IGP SRGB
18,000 – 19,999

```
RP/0/0/CPU0:xrvr-1#show mpls label table detail
Table Label   Owner                              State
----- ------- ---------------------------------- ------
<...snip...>
0     18000   ISIS(A):1                          InUse  No
  (Lbl-blk SRGB, vers:0, (start_label=18000) size=2000)
<...snip...>
```

IS-IS SRGB

Start_label = 18,000

Size = 2,000

Non-default SRGB label block allocation for ISIS
[18,000 – 19,999]

Segment Routing Global Block (SRGB)
Example

```
segment-routing
 global-block 18000 19999
!
router ospf 1
 segment routing mpls
 segment-routing global-block 20000 21999
```

Configure a non-default global SRGB 18,000 – 19,999

Configure an IGP SRGB 20,000 – 21,999

```
RP/0/0/CPU0:xrvr-1#show mpls label table detail
Table Label   Owner                              State
----- ------- ---------------------------------- ------
<...snip...>
0     20000   OSPF(A):ospf-1                     InUse  No
  (Lbl-blk SRGB, vers:0, (start label=20000) size=2000)
<...snip...>
```

OSPF SRGB

Start_label = 20,000

Size = 2,000

Non-default SRGB label block allocation for OSPF
[20,000 – 21,999]

IS-IS SRGB in Router Capability TLV

- The SRGB is advertised in the IS-IS Router Capability TLV
- The Router Capability TLV contains: Router ID (32 bits), Flags (8 bits) and optional sub-TLVs.
- The Router Capability TLV flags:

```
+-+-+-+-+-+-+-+-+
| Reserved  |D|S|
+-+-+-+-+-+-+-+-+
```

  - S: Scope, if set, then flood the TLV across the entire routing domain
  - D: Down, set if the TLV is leaked from level-2 to level-1

IS-IS Segment Routing Capability sub-TLV

- The SR Capabilities sub-TLV is included in the Router Capability TLV
  - The SR Algorithm sub-TLV can also be included. Not in IOS XR 5.3.2
- The SR Capabilities sub-TLV contains: Flags (8 bits) and one or more SRGB descriptors
  - An SRGB descriptor contains: Range (24 bits), SID/Label (variable, 32 bits if MPLS) indicating start of SRGB
- The flags in the SR-Capabilities sub-TLV:

  - I: IPv4, if set, then the router is capable of outgoing IPv4 encapsulation on all interfaces
  - V: IPv6, if set, then the router is capable of outgoing IPv6 encapsulation on all interfaces

IS-IS example

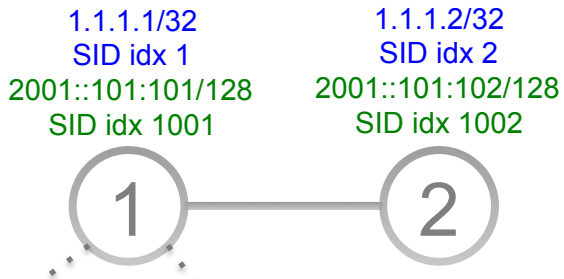

IPv4 elements in blue
IPv6 elements in green

```
router isis 1
 is-type level-2-only
 net 49.0001.0000.0000.0001.00
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 address-family ipv6 unicast
  metric-style wide
  segment-routing mpls
 !
```

```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

IS-IS 1 (Level-2) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
xrvr-1.00-00        * 0x0000039b   0xfc27        1079              0/0/0
  Area Address: 49.0001
  NLPID:        0xcc
  NLPID:        0x8e
  MT:           Standard (IPv4 Unicast)
  MT:           IPv6 Unicast
  Hostname:     xrvr-1
  IP Address:   1.1.1.1
  IPv6 Address: 2001::101:101
  Router Cap:   1.1.1.1, D:0, S:0
    Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000
<...>
```

S:0, don't propagate TLV to other level
D:0, TLV not propagated L2→L1

I:1 IPv4 support
V:1, IPv6 support

SRGB: Start: 16,000, Size: 8000
[16,000 – 23,999]

OSPF SRGB in Router Information LSA

- One or more SID/Label Range TLVs (SRGB descriptors) are included in the Router Information Opaque LSA
  - The SR Algorithm TLV is also included in the Router Info LSA
- The SID/Label Range TLV contains: the Range Size (24 bits) and SID/Label TLV (variable, 32 bits if MPLS) indicating start or SRGB
- The SR Algorithm TLV contains a list of algorithm identifiers (8 bits per identifier) used by the node
  - Algorithm 0: Shortest Path First (SPF) algorithm based on link metric

OSPF example

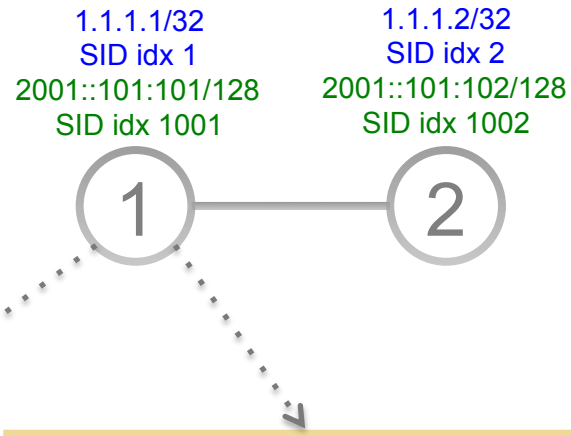


```
router ospf 1
 router-id 1.1.1.1
 segment-routing mpls
 segment-routing forwarding mpls
 area 0
  interface Loopback0
   passive enable
   prefix-sid index 1
  !
  interface GigabitEthernet0/0/0/0
   network point-to-point
  !
 !
!
```

```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 4.0.0.0 self-originate
            OSPF Router with ID (1.1.1.1) (Process ID 1)
                Type-10 Opaque Link Area Link States (Area 0)
<...>
  Opaque Type: 4
<...>
  Advertising Router: 1.1.1.1
<...>
    Router Information TLV: Length: 4
    Capabilities:
      Graceful Restart Helper Capable
      Stub Router Capable
      All capability bits: 0x60000000

    Segment Routing Algorithm TLV: Length: 1
      Algorithm: 0

    Segment Routing Range TLV: Length: 12
      Range Size: 8000

        SID sub-TLV: Length 3
         Label: 16000
```

Only Algorithm 0 (SPF)

SRGB: Start: 16,000, Size: 8000
[16,000 – 23,999]

Prefix-SID
and Adjacency-SID

SID Encoding

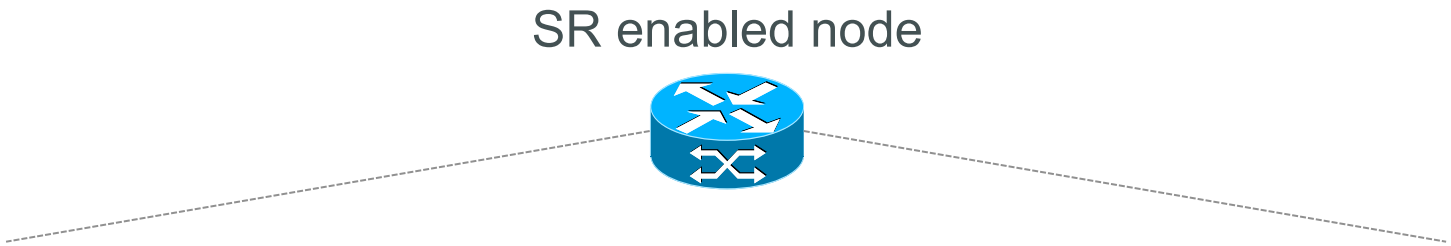

- Prefix SID
  - Uses SR Global Block (SRGB)
  - SRGB advertised with router capabilities TLV
  - In the configuration, Prefix-SID can be configured as an absolute value or an index
  - In the protocol advertisement, Prefix-SID is always encoded as a globally unique index

    Index represents an offset from SRGB base, zero-based numbering, i.e. 0 is 1st index

    E.g. index **1** → SID is 16,000 + **1** = 16,001
- Adjacency SID
  - Locally significant
  - Automatically allocated for each adjacency
  - Always encoded as an absolute (i.e. not indexed) value

Prefix-SID

Prefix segment

- Global Segment – Global significance
  - Unique within SR domain
- Managed by routing protocol
  - IGP allocates a block of labels (SRGB) from Label Switching Database (LSD)
- Manually configured
  - Under IGP enabled loopback interface
  - Only /32 or /128 prefixes in global routing table
- Prefix-SIDs are assigned by the operator similar to e.g. assigning loopback addresses

Node segment

- Node segment is a Prefix segment associated with a host prefix that identifies a node
  - Equivalent to a router-id prefix, which is a prefix identifying a node
  - Node-SID is prefix-SID with N-flag set in advertisement
- By default, each configured prefix-SID is a node-SID
  - “regular” (i.e. non Node-SID) prefix-SID is configurable for IS-IS

Prefix-SID / Node-SID Configuration

```
router isis 1
 interface Loopback0
  address-family ipv4|ipv6 unicast
   prefix-sid {absolute|index} {<SID value>|<SID index>}
```

```
router ospf 1
 area 0
  interface Loopback0
   prefix-sid {absolute|index} {<SID value>|<SID index>}
```

- Prefix-SID can be specified using:
  an absolute value within the SRGB (“global mode”)
  or an index (offset) from the lower bound of the SRGB.

IS-IS example

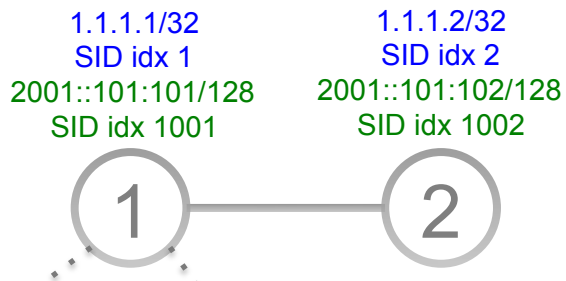

IPv4 elements in blue
IPv6 elements in green

```
interface Loopback0
 ipv4 address 1.1.1.1 255.255.255.255
 ipv6 address 2001::101:101/128
!
router isis 1
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 address-family ipv6 unicast
  metric-style wide
  segment-routing mpls
 !
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid absolute 16001
   !! Or: prefix-sid index 1
  !
  address-family ipv6 unicast
   prefix-sid absolute 17001
   !! Or: prefix-sid index 1001
 !
```

```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

IS-IS 1 (Level-2) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
xrvr-1.00-00        * 0x0000039b   0xfc27        1079            0/0/0
  Area Address: 49.0001
  NLPID:        0xcc
  NLPID:        0x8e
  MT:           Standard (IPv4 Unicast)
  MT:           IPv6 Unicast                                     0/0/0
  Hostname:     xrvr-1
  IP Address:   1.1.1.1
  IPv6 Address: 2001::101:101
  Router Cap:   1.1.1.1, D:0, S:0
    Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000
<...>
  Metric: 0          IP-Extended 1.1.1.1/32
    Prefix-SID Index: 1, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
<...>
  Metric: 0          MT (IPv6 Unicast) IPv6 2001::101:101/128
    Prefix-SID Index: 1001, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
<...>
```

IS-IS Prefix-SID flags

- IS-IS Prefix-SID advertisements have the following flags:

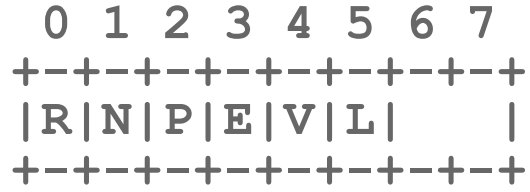

  - R: Re-advertisement, set if the attached non-local prefix is propagated to another level or redistributed – default in IOS XR: 0
  - N: Node-SID, set if the prefix-SID is a node-SID, i.e. identifies the node – default in IOS XR: 1
  - P: no-PHP, set if the penultimate hop must NOT pop the prefix-SID before forwarding the packet – default in IOS XR: 0
  - E: Explicit-Null, set if penultimate hop must replace prefix-SID with Explicit-Null label – default in IOS XR: 0
  - V: Value, set if prefix-SID carries a value (not an index) – IOS XR: always unset
  - L: Local, set if prefix-SID has local significance – IOS XR: always unset

OSPF Extended Prefix TLV flags

- OSPF Extended Prefix TLVs have the following flags:

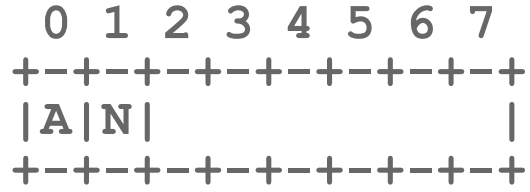

  - A: Attach, set if ABR generates an Extended Prefix TLV for inter-area prefix that is locally connected or attached in another connected area – default in IOS XR: 0
  - N: Node, set if the prefix identifies the node – default in IOS XR: 1

OSPF Prefix-SID sub-TLV flags

- OSPF Prefix-SID sub-TLVs have the following flags:

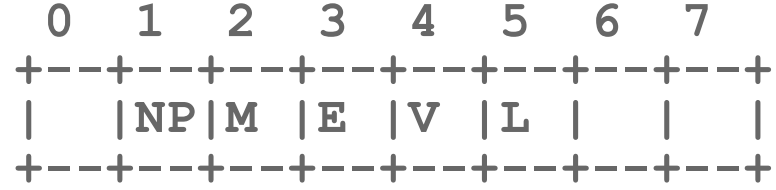

  - NP: no-PHP, set if the penultimate hop must NOT pop the prefix-SID before forwarding the packet – default in IOS XR: 0
  - M: Mapping Server, set if the SID is advertised from the Mapping Server functionality – default in IOS XR: 0
  - E: Explicit-Null, set if penultimate hop must replace prefix-SID with Explicit-Null label – default in IOS XR: 0
  - V: Value, set if prefix-SID carries a value (not an index) – IOS XR: always unset
  - L: Local, set if prefix-SID has local significance – IOS XR: always unset

Node-SID flag (N-flag)

- By default, the N-flag is set on each configured Prefix-SID
  - The Prefix-SID is then also a Node-SID
- To clear the N-flag, configure `n-flag-clear`
  - The N-flag should be cleared for e.g. Anycast-SIDs

```
prefix-sid absolute 18000 n-flag-clear
```

- Applies to both IS-IS and OSPF

IS-IS Node-SID flag (N-flag) – Example

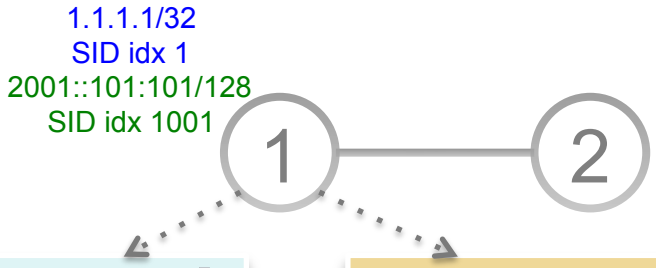

IPv4 elements in blue
IPv6 elements in green

```
interface Loopback0
 ipv4 address 1.1.1.1 255.255.255.255
 ipv6 address 2001::101:101/128
!
router isis 1
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid absolute 16001
  !
  address-family ipv6 unicast
   prefix-sid absolute 17001
```

```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

<...>
  Metric: 0           IP-Extended 1.1.1.1/32
    Prefix-SID Index: 1, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
<...>
  Metric: 0           MT (IPv6 Unicast) IPv6 2001::101:101/128
    Prefix-SID Index: 1001, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
<...>
```

```
interface Loopback0
 ipv4 address 1.1.1.1 255.255.255.255
 ipv6 address 2001::101:101/128
!
router isis 1
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid absolute 16001 n-flag-clear
  !
  address-family ipv6 unicast
   prefix-sid absolute 17001 n-flag-clear
```

```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

<...>
  Metric: 0           IP-Extended 1.1.1.1/32
    Prefix-SID Index: 1, Algorithm:0, R:0 N:0 P:0 E:0 V:0 L:0
<...>
  Metric: 0           MT (IPv6 Unicast) IPv6 2001::101:101/128
    Prefix-SID Index: 1001, Algorithm:0, R:0 N:0 P:0 E:0 V:0 L:0
<...>
```

Prefix-SID – explicit-Null and no-PHP Flags

- If the originator of a Prefix-SID wishes to receive packets with an MPLS label to enable using the MPLS EXP/TC bits for classification
  - The penultimate hop can swap the top label with explicit-null label instead of popping it, then the packets arrive with a label, with MPLS EXP/TC bits
- To request the explicit-null behavior, the originator of the Prefix-SID can set the E-flag in the Prefix-SID advertisement
  - By default, the E-flag is unset

Prefix-SID – explicit-Null and no-PHP Flags

- To advertise prefix-SID with the E-flag set (requesting explicit-null behavior), use the **`explicit-null`** keyword on the prefix-SID configuration

```
prefix-sid absolute 16001 explicit-null
```

  - This automatically sets no-PHP flag ((N)P-flag*) as well (no-PHP cannot be configured independently in IOS XR)
- Applies to both IS-IS and OSPF

*This flag is named P-flag in IS-IS IETF draft, NP-flag in OSPF IETF draft

IS-IS E-flag and (N)P-Flag – Example

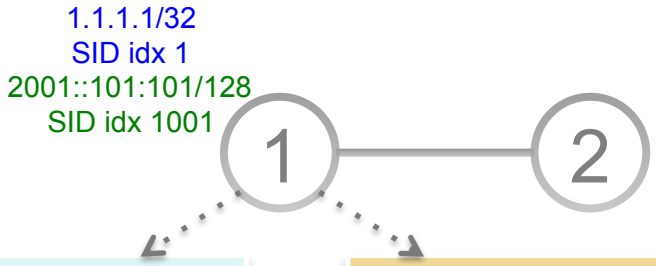

IPv4 elements in blue
IPv6 elements in green

```
interface Loopback0
 ipv4 address 1.1.1.1 255.255.255.255
 ipv6 address 2001::101:101/128
!
router isis 1
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid absolute 16001
  !
  address-family ipv6 unicast
   prefix-sid absolute 17001
```

```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

<...>
  Metric: 0           IP-Extended 1.1.1.1/32
    Prefix-SID Index: 1, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
<...>
  Metric: 0           MT (IPv6 Unicast) IPv6 2001::101:101/128
    Prefix-SID Index: 1001, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
<...>
```

```
interface Loopback0
 ipv4 address 1.1.1.1 255.255.255.255
 ipv6 address 2001::101:101/128
!
router isis 1
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid absolute 16001 explicit-null
  !
  address-family ipv6 unicast
   prefix-sid absolute 17001 explicit-null
```

```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

<...>
  Metric: 0           IP-Extended 1.1.1.1/32
    Prefix-SID Index: 1, Algorithm:0, R:0 N:1 P:1 E:1 V:0 L:0
<...>
  Metric: 0           MT (IPv6 Unicast) IPv6 2001::101:101/128
    Prefix-SID Index: 1001, Algorithm:0, R:0 N:1 P:1 E:1 V:0 L:0
<...>
```

OSPF E-flag and (N)P-Flag – Example

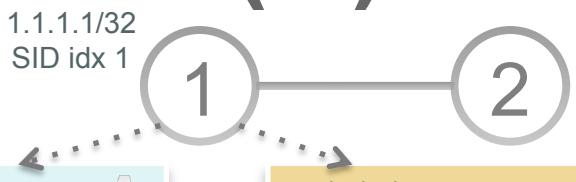


```
interface Loopback0
 ipv4 address 1.1.1.1 255.255.255.255
!
router ospf 1
 area 0
  interface Loopback0
   prefix-sid index 1
```

```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 7.0.0.1 self-originate
 OSPF Router with ID (1.1.1.1) (Process ID 1)
                Type-10 Opaque Link Area Link States (Area 0)
<...>
    Extended Prefix TLV: Length: 20
      Route-type: 1
      AF        : 0
      Flags     : 0x40              0  1  2  3  4  5  6  7
      Prefix    : 1.1.1.1/32      +--+--+--+--+--+--+--+--+
                                  |  |NP|M |E |V |L |  |  |
      SID sub-TLV: Length: 8      +--+--+--+--+--+--+--+--+
        Flags     : 0x0  ------>  0  0  0  0  0  0  0  0  -> all flag unset
        MTID      : 0
        Algo      : 0
        SID Index : 1
```

```
interface Loopback0
 ipv4 address 1.1.1.1 255.255.255.255
!
router ospf 1
 area0
  interface Loopback0
   prefix-sid index 1 explicit-null
```

```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 7.0.0.1 self-originate
 OSPF Router with ID (1.1.1.1) (Process ID 1)
                Type-10 Opaque Link Area Link States (Area 0)
<...>
    Extended Prefix TLV: Length: 20
      Route-type: 1
      AF        : 0
      Flags     : 0x40              0  1  2  3  4  5  6  7
      Prefix    : 1.1.1.1/32      +--+--+--+--+--+--+--+--+
                                  |  |NP|M |E |V |L |  |  |
      SID sub-TLV: Length: 8      +--+--+--+--+--+--+--+--+
        Flags     : 0x50  ------>  0  1  0  1  0  0  0  0  -> NP-flag and E-flag set
        MTID      : 0
        Algo      : 0
        SID Index : 1
```

Anycast Prefix Segments

- Anycast prefixes: same prefix advertised by multiple nodes
- Anycast prefix-SID: prefix-SID associated with anycast prefix
  - Same prefix-SID for the same prefix!
- Traffic is forwarded to one of the Anycast prefix-SID originators based on best IGP path
- If primary node fails, traffic is auto re-routed to the other node
- Note: nodes advertising the same Anycast prefix-SID must have the same SRGB

Anycast Prefix Segments – Benefits

- Coarse Grained Traffic Engineering, steering traffic via groups of routers (with common Anycast-SID)
  - latency policy example:
    - > from Japan to Western Europe: SPT is via US. For low-latency: steer traffic on SPT to Eastern Europe (SPT Japan→Eastern Europe is low-latency), then on SPT to destination in Western Europe (SPT Eastern Europe→Western Europe is low-latency).
    - > The intermediate hop can be any BB router in Eastern Europe: steer traffic via “Eastern Europe BB router” Anycast-SID from Japan to Eastern Europe
- High-availability
  - if one of the Eastern routers fail, the policy survives
- Typical for service virtualization
  - nearest firewall/DPI etc.

Anycast-SID – High Availability benefit

- PE1 and PE2 each advertise a prefix-SID, 16001 resp. 16002
- PE1 and PE2 both advertise an Anycast prefix-SID, 16100

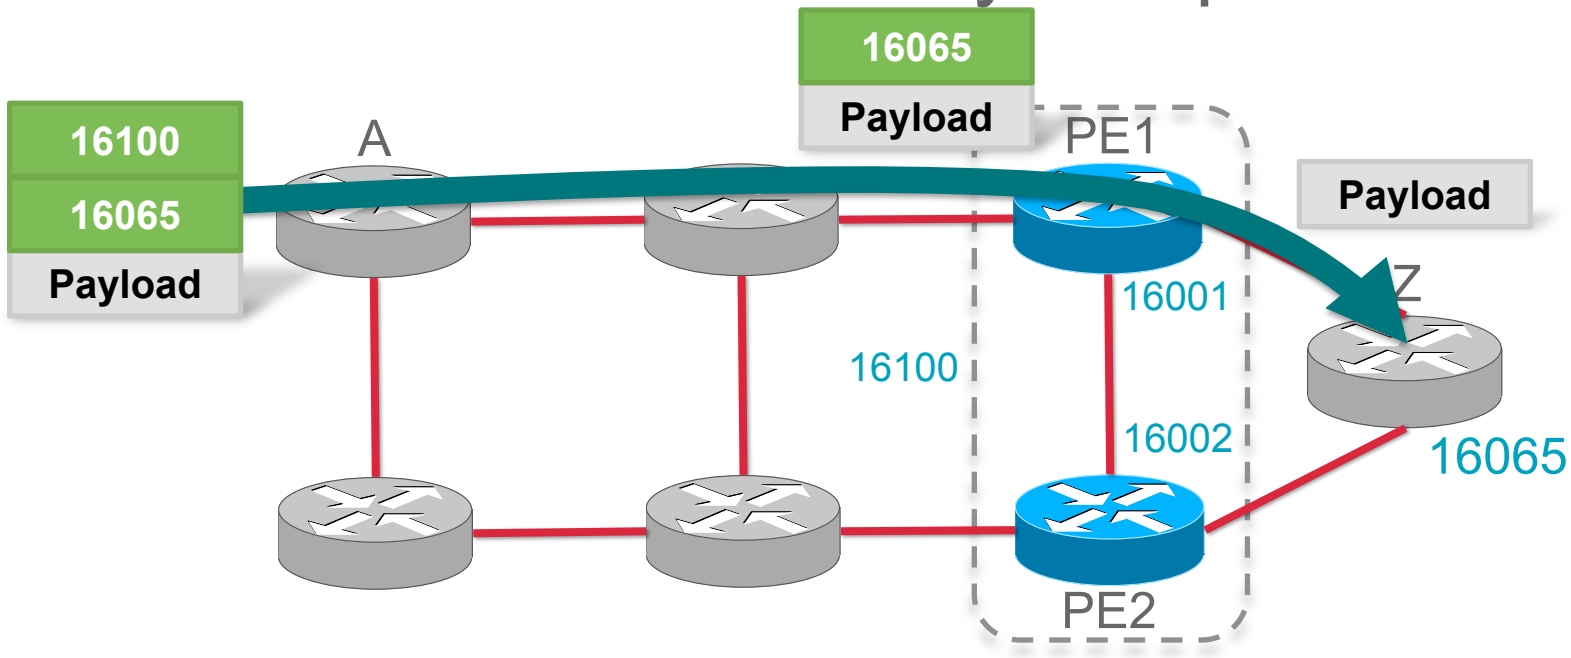

Anycast-SID – High Availability benefit

- If closest anycast node fails, traffic is auto re-routed to another node advertising the Anycast prefix-SID

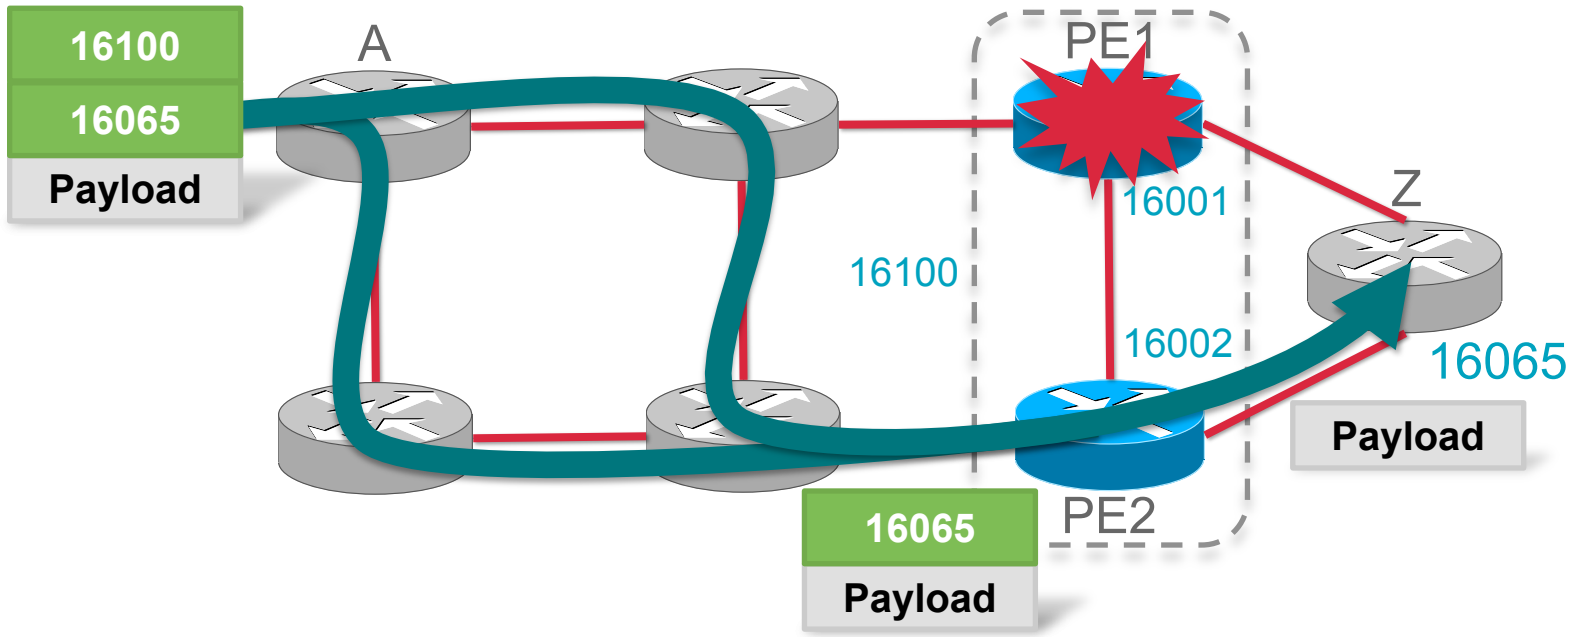

Adjacency-SID

Adjacency segments

- Local segment – Local significance
  - Local label, allocated from dynamic label pool
- Automatically allocated for each adjacency
  - Per adjacency: a protected and an unprotected adjacency-SID
    - >See SRTE presentation for more information
  - IS-IS: Different Adjacency-SID for L1 and L2 adjacencies between same neighbors
  - IS-IS: Different Adjacency-SID for IPv4 and IPv6 address-families
  - OSPF: Same Adjacency-SID in all areas of Multi-Area Adjacency (multiple adjacencies, each for a different area, over same interface)

Adjacency-SID label

```
RP/0/0/CPU0:xrvr-1#show ospf neighbor 1.1.1.2 detail
Neighbors for OSPF 1

 Neighbor 1.1.1.2, interface address 99.1.2.2
    In the area 0 via interface GigabitEthernet0/0/0/0
    Neighbor priority is 1, State is FULL, 6 state changes
<...>
    Adjacency SID Label: 24006
```

```
RP/0/0/CPU0:xrvr-1#show isis adjacency detail
IS-IS 1 Level-2 adjacencies:
System Id      Interface        SNPA           State Hold Changed  NSF IPv4 IPv6
                                                                     BFD  BFD
xrvr-2         Gi0/0/0/0        *PtoP*         Up    24   1w1d     Yes None None
  Area Address:            49.0001
  Neighbor IPv4 Address:   99.1.2.2*
  Adjacency SID:           24000
  Neighbor IPv6 Address:   fe80::f816:3eff:febd:b683*
  Adjacency SID:           24003
  Topology:                IPv4 Unicast
  Topology:                IPv6 Unicast
```

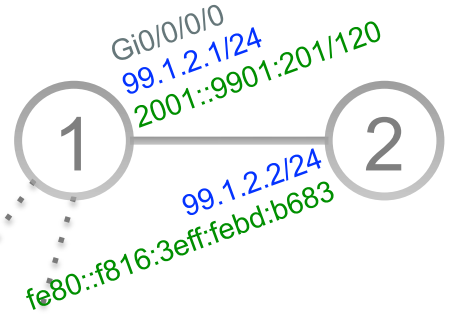

Adjacency SID label

IPv4 Adjacency SID label

IPv6 Adjacency SID label

IS-IS Adjacency-SID sub-TLV flags

- IS-IS Adjacency-SID advertisements have the following flags:

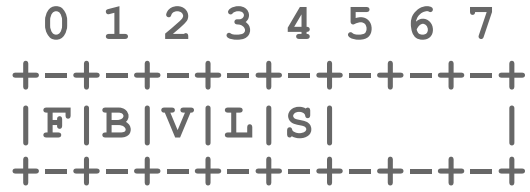

  - F: Address-Family, unset: IPv4, set: IPv6
  - B: Backup, set if the Adj-SID refers to a protected adjacency (e.g. Using TI-LFA)
  - V: Value, set if Adj-SID carries a value – IOS XR: always set
  - L: Local, set if Adj-SID has local significance – IOS XR: always set
  - S: Set, set if Adj-SID refers to a set of adjacencies – IOS XR: always unset
- IS-IS Adjacency-SID sub-TLV contains a Weight field (8 bits): The value represents the weight of the Adj-SID for the purpose of load balancing
  - Weight = 0 in IOS XR 5.3.2

OSPF Adacency-SID sub-TLV flags

- OSPF Adjacency-SID sub-TLVs have the following flags:

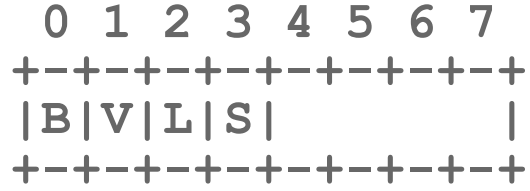

  - B: Backup, set if the Adj-SID refers to a protected adjacency (e.g. Using TI-LFA)
  - V: Value, set if Adj-SID carries a value (not an index) – IOS XR: always set
  - L: Local, set if Adj-SID has local significance – IOS XR: always set
  - S: Set, set if Adj-SID refers to a set of adjacencies – IOS XR: always unset
- OSPF Adjacency-SID sub-TLV contains a Weight field (8 bits): The value represents the weight of the Adj-SID for the purpose of load balancing
  - Weight = 0 in IOS XR 5.3.2

Adjacency segments – IS-IS Example

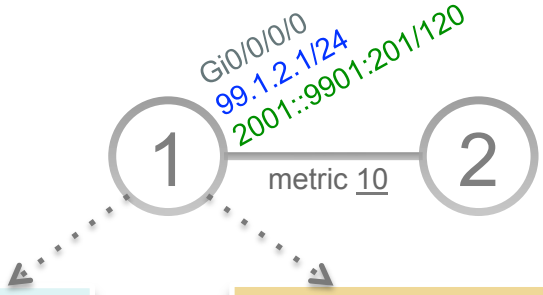

IPv4 elements in blue
IPv6 elements in green

```
interface GigabitEthernet0/0/0/0
 ipv4 address 99.1.2.1 255.255.255.0
 ipv6 address 2001::9901:201/120
!
router isis 1
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 address-family ipv6 unicast
  metric-style wide
  segment-routing mpls
 !
 interface GigabitEthernet0/0/0/0
  point-to-point
  address-family ipv4 unicast
  address-family ipv6 unicast
```

```
RP/0/0/CPU0:xrvr-1#show isis adjacency detail

IS-IS 1 Level-2 adjacencies:
System Id      Interface        SNPA           State Hold Changed  NSF IPv4 IPv6
                                                                      BFD  BFD
xrvr-2         Gi0/0/0/0        *PtoP*         Up    24   1w1d     Yes None None
  Area Address:           49.0001
  Neighbor IPv4 Address:  99.1.2.2*
  Adjacency SID:          24000
  Neighbor IPv6 Address:  fe80::f816:3eff:febd:b683*
  Adjacency SID:          24003
  Topology:               IPv4 Unicast
  Topology:               IPv6 Unicast

Total adjacency count: 1
```

Adjacency segments – IS-IS Example

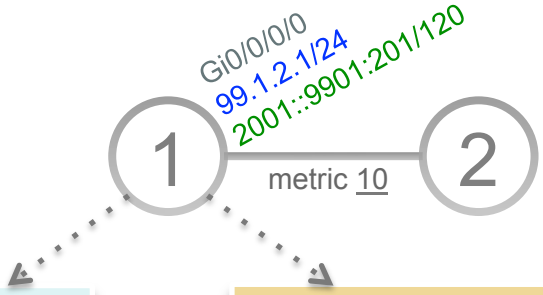

IPv4 elements in blue
IPv6 elements in green

```
interface GigabitEthernet0/0/0/0
 ipv4 address 99.1.2.1 255.255.255.0
 ipv6 address 2001::9901:201/120
!
router isis 1
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 address-family ipv6 unicast
  metric-style wide
  segment-routing mpls
 !
 interface GigabitEthernet0/0/0/0
  point-to-point
  address-family ipv4 unicast
  address-family ipv6 unicast
```

```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

IS-IS 1 (Level-2) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
xrvr-1.00-00        * 0x0000039b   0xfc27        1079            0/0/0
  Area Address: 49.0001
  NLPID:        0xcc
  NLPID:        0x8e
  MT:           Standard (IPv4 Unicast)
  MT:           IPv6 Unicast                                     0/0/0
  Hostname:     xrvr-1
  IP Address:   1.1.1.1
  IPv6 Address: 2001::101:101
  Router Cap:   1.1.1.1, D:0, S:0
    Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000
  Metric: 10         IS-Extended xrvr-2.00
    Interface IP Address: 99.1.2.1
    Neighbor IP Address: 99.1.2.2
    ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24000
  Metric: 10         IP-Extended 99.1.2.0/24
<...>
  Metric: 10         MT (IPv6 Unicast) IS-Extended xrvr-2.00
    ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24003
  Metric: 10         MT (IPv6 Unicast) IPv6 2001::9901:200/120
<...>
```

Adjacency segments – OSPF Example

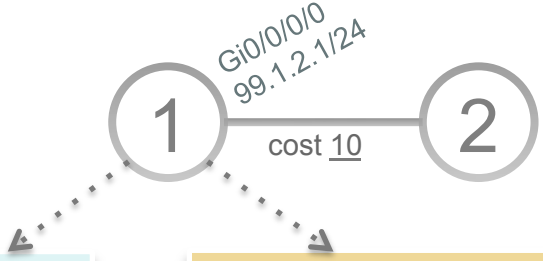


```
interface GigabitEthernet0/0/0/0
 ipv4 address 99.1.2.1 255.255.255.0
!
router ospf 1
 segment-routing mpls
 segment-routing forwarding mpls
 area 0
  interface GigabitEthernet0/0/0/0
  network point-to-point
```

```
RP/0/0/CPU0:xrvr-1#show ospf neighbor detail
Neighbors for OSPF 1

 Neighbor 1.1.1.2, interface address 99.1.2.2
    In the area 0 via interface GigabitEthernet0/0/0/0
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 0.0.0.0 BDR is 0.0.0.0
    Options is 0x52
    LLS Options is 0x1 (LR)
    Dead timer due in 00:00:36
    Neighbor is up for 00:00:31
    Number of DBD retrans during last exchange 0
    Index 1/1, retransmission queue length 0, number of retransmission 1
    First 0(0)/0(0) Next 0(0)/0(0)
    Last retransmission scan length is 1, maximum is 1
    Last retransmission scan time is 0 msec, maximum is 0 msec
    LS Ack list: NSR-sync pending 0, high water mark 0
    Adjacency SID Label: 24006
```

Adjacency segments – OSPF Example

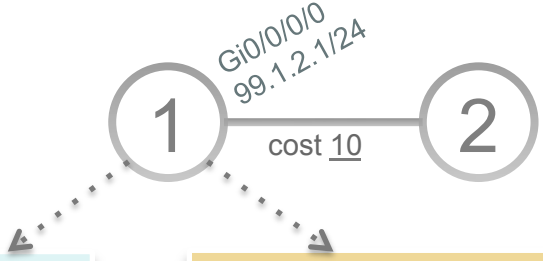


```
interface GigabitEthernet0/0/0/0
 ipv4 address 99.1.2.1 255.255.255.0
!
router ospf 1
 segment-routing mpls
 segment-routing forwarding mpls
 area 0
  interface GigabitEthernet0/0/0/0
  network point-to-point
```

```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 8.0.0.1 self-originat
            OSPF Router with ID (1.1.1.1) (Process ID 1)
                Type-10 Opaque Link Area Link States (Area 0)
<...>
  Opaque Type: 8
  Opaque ID: 1
  Advertising Router: 1.1.1.1
  LS Seq Number: 80000029
  Checksum: 0x9147
  Length: 48

    Extended Link TLV: Length: 24
      Link-type : 1
      Link ID   : 1.1.1.2
      Link Data : 99.1.2.1

     Adj sub-TLV: Length: 7
        Flags     : 0x60
        MTID      : 0
        Weight    : 0
        Label     : 24006
```

LAN Adjacency-SID

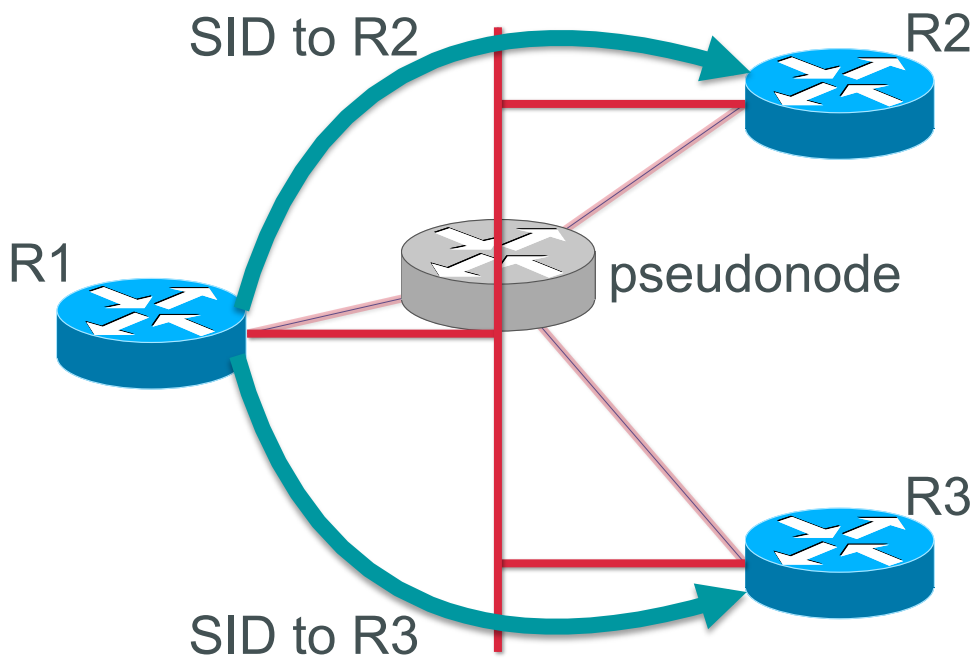

- All nodes on a LAN advertise their adjacency to a pseudonode only
  - The pseudonode represents the network
- For SR to steer traffic to each node on the LAN, an Adjacency-SID is needed to each other node on the LAN
  - These LAN-Adj-SIDs are associated with the adjacency to pseudonode
- E.g. Node R1 will allocate and advertise a LAN Adj-SID to node R2 and one to node R3

OSPF LAN Adjacency-SID

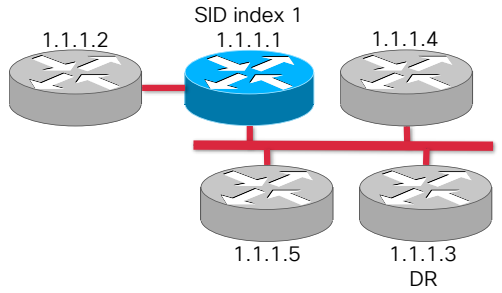


```
RP/0/0/CPU0:xrvr-1#show ospf neighbor detail | i "interface addres
 Neighbor 1.1.1.4, interface address 66.0.0.4
    Neighbor priority is 1, State is FULL, 6 state changes
    Adjacency SID Label: 24002
 Neighbor 1.1.1.5, interface address 66.0.0.5
    Neighbor priority is 1, State is 2WAY, 2 state changes
    Adjacency SID Label: 24009
 Neighbor 1.1.1.3, interface address 66.0.0.3
    Neighbor priority is 1, State is FULL, 6 state changes
    Adjacency SID Label: 24000
```

Adjacencies to LAN nodes

Adjacency to BDR node

Adjacency to non-(B)DR node

Adjacency to DR node

- OSPF allocates an Adjacency-SID for each adjacency in the 2WAY state or higher
  - Reminder: On a broadcast or NMBA network, adjacencies to DR and BDR reach FULL state. Adjacencies to other nodes on the network (DR-OTHER) reach 2WAY state.
- On a broadcast or NMBA network, a node advertises an Adj-SID using the Adj-SID sub-TLV for its adjacency to the DR and advertises Adj-SIDs using the LAN Adj-SID for other neighbors (e.g. BDR, DR-OTHER) on the network

OSPF LAN Adj-SID advertisement

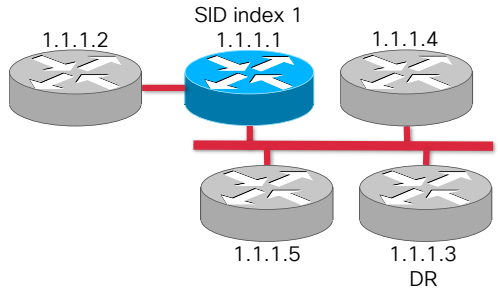


```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 8.0.0.7 self-originate
<...>
 Opaque Type: 8
<...>
 Extended Link TLV: Length: 40
      Link-type : 2
      Link ID   : 66.0.0.3
      Link Data : 66.0.0.1

     LAN Adj sub-TLV: Length: 11
       Flags       : 0x60
       MTID        : 0
       Weight      : 0
       Neighbor ID: 1.1.1.4
       Label       : 24002

     LAN Adj sub-TLV: Length: 11
       Flags       : 0x60
       MTID        : 0
       Weight      : 0
       Neighbor ID: 1.1.1.5
       Label       : 24009

     Adj sub-TLV: Length: 7
        Flags      : 0x60
        MTID       : 0
        Weight     : 0
        Label      : 24000
```

Opaque type 8 (Extended Link)

Adjacency: link to transit network

LAN Adjacency-SID for neighbor 1.1.1.4: label 24002

LAN Adjacency-SID for neighbor 1.1.1.5: label 24009

Adjacency-SID to Designated Router: label 24000

IS-IS LAN Adj-SID advertisement

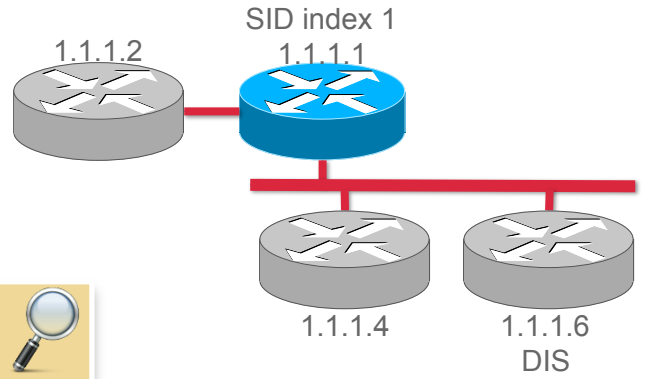


```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1
<...>
  Metric: 10         IS-Extended xrvr-6.01
    Interface IP Address: 99.1.6.1
    Neighbor IP Address: 99.1.6.6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24009 System ID:xrvr-6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24007 System ID:xrvr-4
<...>
```

- IS-IS attaches the LAN Adj-SID sub-TLVs to the TLV of the adjacency to the pseudonode (DIS)
  - One LAN Adj-SID sub-TLV for each neighbor on the LAN

LAN Adjacency-SID label

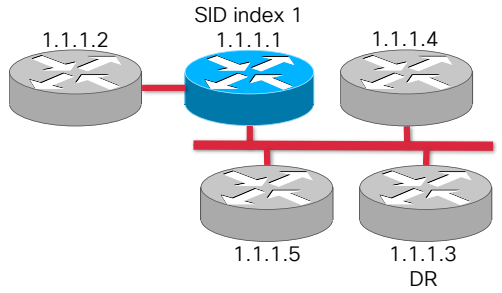


```
RP/0/0/CPU0:xrvr-1#show mpls forwarding | i "L|--|24002|24009|24000"
Local  Outgoing    Prefix             Outgoing     Next Hop        Bytes
Label  Label       or ID              Interface                    Switched
------ ----------- ------------------ ------------ --------------- ------------
24000  Pop         SR Adj (idx 0)     Gi0/0/0/3    66.0.0.3        0
24002  Pop         SR Adj (idx 0)     Gi0/0/0/3    66.0.0.4        0
24009  Pop         SR Adj (idx 0)     Gi0/0/0/3    66.0.0.5        0
```

Same outgoing interface

Different next hop

OPSF multi-area
IS-IS multi-level

Terminology

- "advertise":
  - When a node advertises a prefix, it includes that prefix in the link state advertisements it generates and sends to its neighbors
- "originate":
  - When a node originates a prefix, it advertises a local prefix, a prefix "owned" by the node
- "propagate" (or "re-advertise"):
  - When a node propagates a prefix, it advertises a prefix in an area or level that it has received from another area or level, or that it has originated in another area or level
    - In this presentation, the node is an ABR or L1L2 router
  - A propagated prefix can be local or non-local

Multi-area, multi-level

- Enabling Segment Routing does not change how multi-area, multi-level works
- Prefix-SIDs are propagated between areas
  - Attached/associated with the propagated prefix
  - Some prefix-SID flags are modified when propagated
- Adjacency-SIDs are not propagated between areas

Propagating Prefix-SID – flags

- When Area Border Router or L1L2 router propagates a non-local prefix with prefix-SID
  - Set ”PHP-off” flag** of prefix-SID
  - Clear ”Explicit-null” flag of prefix-SID
  - IS-IS: Set “Re-advertisement” flag of prefix-SID
- When Area Border Router or L1L2 router propagates a local prefix with prefix-SID
  - “PHP-off” flag** as configured for the prefix-SID (default: PHP on)
  - “Explicit-null” flag as configured for the prefix-SID (default: no Exp-null)
  - IS-IS: Clear “Re-advertisement” flag of prefix-SID

* “propagated local prefix”: local prefix propagated to another area/level
** “PHP-off” flag is called P-flag in IS-IS draft, NP-flag in OSPF draft

Multi-area OSPF – Example

- R2 propagates remote prefixes (e.g. R1, R3 loopbacks) and its local prefixes (e.g. R2 loopbacks in area0 or area1)
- All prefix-SIDs are configured for PHP (default)

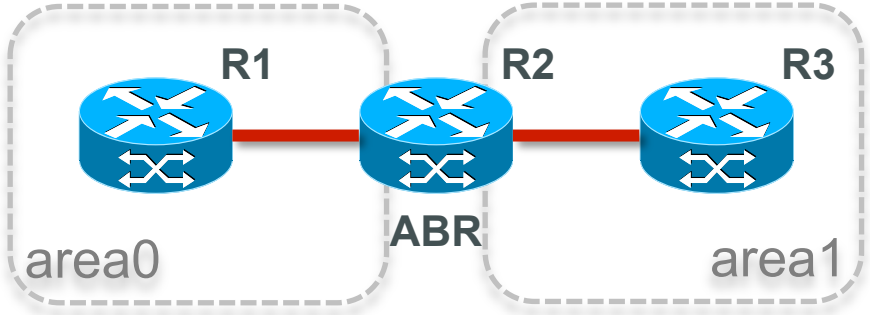

R1: 1.1.1.1/32; SID: 16001 (area0)
R2: 1.1.1.2/32; SID: 16002 (area0)
2.2.2.2/32; SID: 17002 (area1)
R3: 1.1.1.3/32; SID: 16003 (area1)

	Flags in Area 0		Flags in Area 1
1.1.1.1/32; SID 16001	**NP=0**; E=0	→	**NP=1**; E=0
1.1.1.2/32; SID 16002	NP=0; E=0	→	NP=0; E=0
2.2.2.2/32; SID 17002	NP=0; E=0	←	NP=0; E=0
1.1.1.3/32; SID 16003	**NP=1**; E=0	←	**NP=0**; E=0

Propagated local prefixes

Multi-area OSPF – Example

	Flags in Area 0		Flags in Area 1
1.1.1.1/32; SID 16001	**NP=0**; E=0	→	**NP=1**; E=0
1.1.1.2/32; SID 16002	NP=0; E=0	→	NP=0; E=0
2.2.2.2/32; SID 17002	NP=0; E=0	←	NP=0; E=0
1.1.1.3/32; SID 16003	**NP=1**; E=0	←	**NP=0**; E=0

Propagated local prefixes

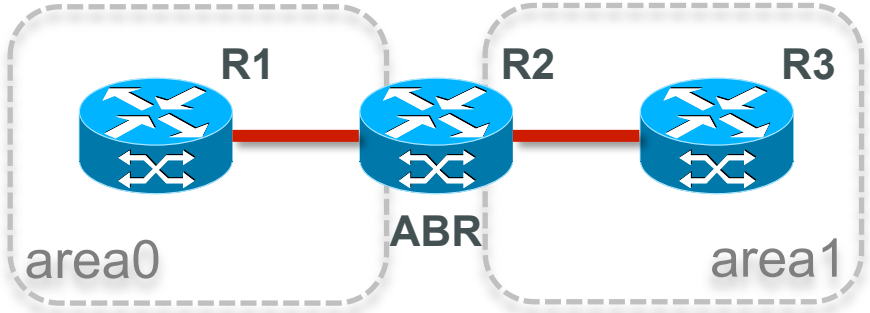

R1: 1.1.1.1/32; SID: 16001 (area0)
R2: 1.1.1.2/32; SID: 16002 (area0)
2.2.2.2/32; SID: 17002 (area1)
R3: 1.1.1.3/32; SID: 16003 (area1)

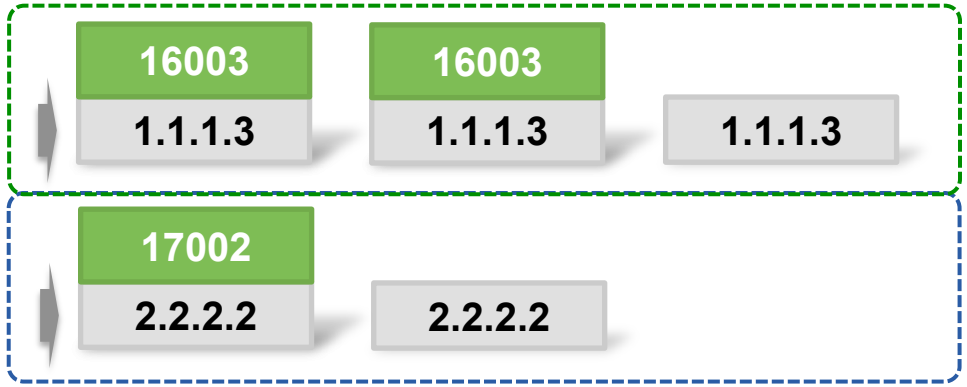

IS-IS Segment Routing Advertisements

IS-IS Configuration – Example

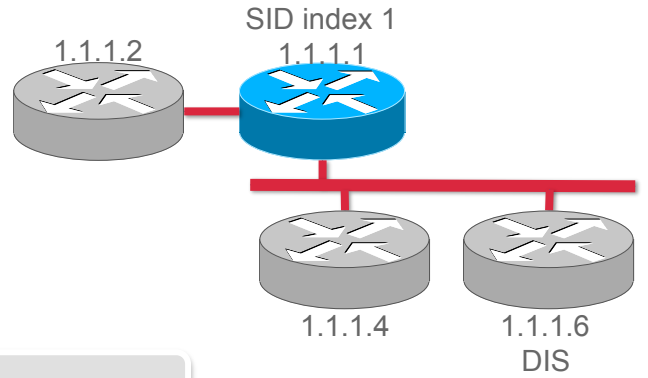


```
router isis 1
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 address-family ipv6 unicast
  metric-style wide
  segment-routing mpls
 !
 interface Loopback0
  passive
  address-family ipv4 unicast
   prefix-sid absolute 16001
  !
  address-family ipv6 unicast
   prefix-sid absolute 20001
  !
 !
<continue...>
```

Wide metrics

enable SR IPv4 control plane and SR MPLS data plane on all ipv4 interfaces in this IS-IS instance

Wide metrics

enable SR IPv6 control plane and SR MPLS data plane on all ipv6 interfaces in this IS-IS instance

Ipv4 Prefix-SID value for loopback0

Ipv6 Prefix-SID value for loopback0

IS-IS Configuration – Example

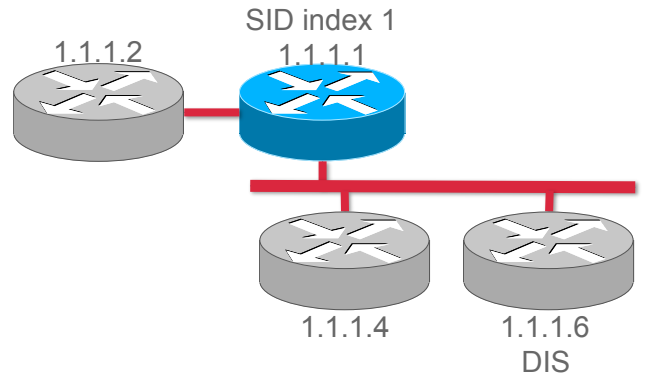


```
<...continue>
 !
 interface TenGigE0/0/2/0
  point-to-point
  address-family ipv4 unicast
  !
  address-family ipv6 unicast
 !
 interface TenGigE0/0/3/0
  address-family ipv4 unicast
  !
  address-family ipv6 unicast
```

Adjacency-SIDs will automatically be allocated for all adjacencies

Verifying IS-IS Database

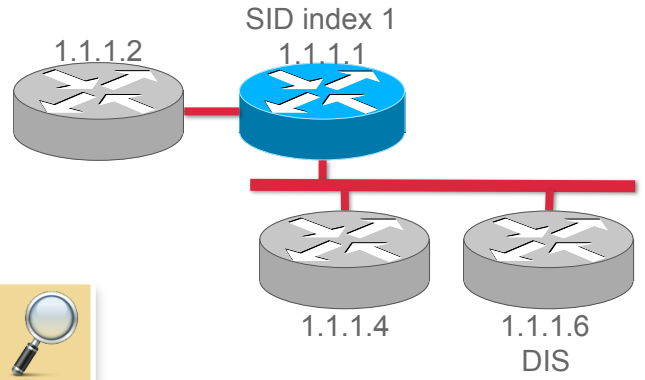


```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

IS-IS 1 (Level-2) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
xrvr-1.00-00        * 0x00000073   0x4eba        1161              0/0/0
  Area Address: 49.0001
  NLPID:        0xcc
  NLPID:        0x8e
  MT:           Standard (IPv4 Unicast)
  MT:           IPv6 Unicast                                     0/0/0
  Hostname:     xrvr-1
  IP Address:   1.1.1.1
  IPv6 Address: 2001::1:1:1:1
  Router Cap:   1.1.1.1, D:0, S:0
    Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000

<...>
```

Two topologies: IPv4 and IPv6

S:0, don't propagate TLV to other level
D:0, TLV not propagated L2→L1

Router capabilities TLV

Verifying IS-IS Database

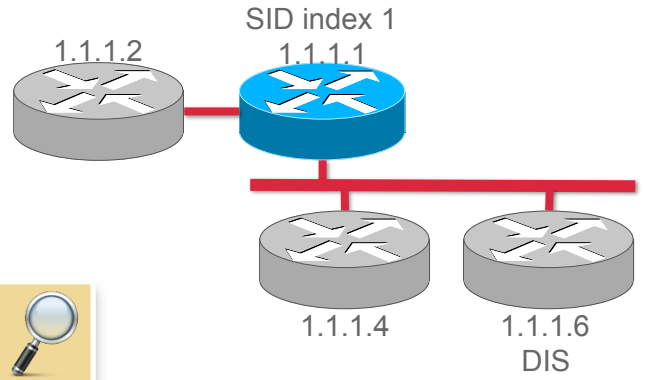


```
RP/0/0/CPU0:xrvr-1#show isis database verbose xrvr-1

IS-IS 1 (Level-2) Link State Database
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
xrvr-1.00-00        * 0x00000073   0x4eba        1161              0/0/0
  Area Address: 49.0001
  NLPID:        0xcc
  NLPID:        0x8e
  MT:           Standard (IPv4 Unicast)
  MT:           IPv6 Unicast                                     0/0/0
  Hostname:     xrvr-1
  IP Address:   1.1.1.1
  IPv6 Address: 2001::1:1:1:1
  Router Cap:   1.1.1.1, D:0, S:0
    Segment Routing: I:1 V:1, SRGB Base: 16000 Range: 8000

<...>
```

Two topologies: IPv4 and IPv6

Single SRGB for both IPv4 and IPv6
I:1 IPv4; V:1 IPv6
Label range 16,000 – 23,999

Verifying IS-IS Database

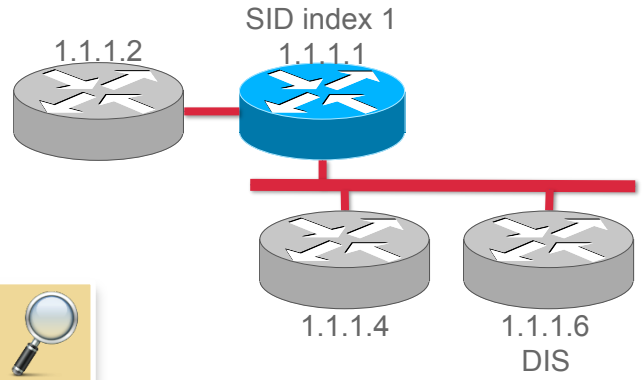


```
<...>

  Metric: 0           IP-Extended 1.1.1.1/32
    Prefix-SID Index: 1, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
  Metric: 10          IP-Extended 99.1.2.0/24
  Metric: 10          IP-Extended 99.1.6.0/24
  Metric: 0           MT (IPv6 Unicast) IPv6 2001::1:1:1:1/128
    Prefix-SID Index: 4001, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
  Metric: 10          MT (IPv6 Unicast) IPv6 2001::99:1:2:0/120
  Metric: 10          MT (IPv6 Unicast) IPv6 2001::99:1:6:0/120

<...>
```

IPv4 IP Reachability TLV

IPv6 IP Reachability TLV

Verifying IS-IS Database

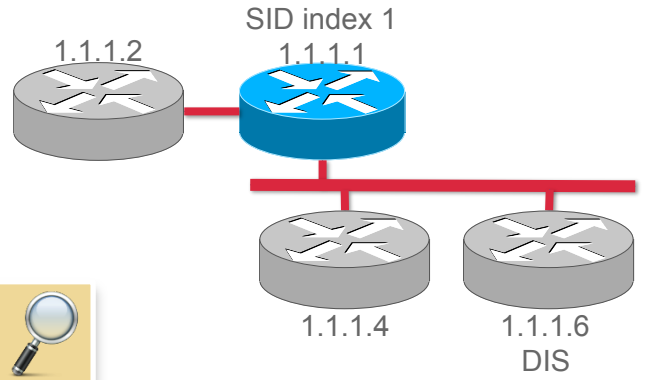


```
<...>

  Metric: 0           IP-Extended 1.1.1.1/32
    Prefix-SID Index: 1, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
  Metric: 10          IP-Extended 99.1.2.0/24
  Metric: 10          IP-Extended 99.1.6.0/24
  Metric: 0           MT (IPv6 Unicast) IPv6 2001::1:1:1:1/128
    Prefix-SID Index: 4001, Algorithm:0, R:0 N:1 P:0 E:0 V:0 L:0
  Metric: 10          MT (IPv6 Unicast) IPv6 2001::99:1:2:0/120
  Metric: 10          MT (IPv6 Unicast) IPv6 2001::99:1:6:0/120

<...>
```

IPv4 prefix-SID
R:0 not propagated
N:1 node SID
P:0 PHP on
E:0 no explicit-Null label

IPv6 prefix-SID
R:0 not propagated
N:1 node SID
P:0 PHP on
E:0 no explicit-Null label

Prefix-SID Flags (see "Prefix Segment" slides):

- R: Re-advertisement, set if the attached non-local prefix is propagated to another level or redistributed – default in IOS XR: 0
- N: Node-SID, set if the prefix-SID is a node-SID, i.e. identifies the node – default in IOS XR: 1
- P: no-PHP, set if the penultimate hop must NOT pop the prefix-SID before forwarding the packet – default in IOS XR: 0
- E: Explicit-Null, set if penultimate hop must replace prefix-SID with Explicit-Null label – default in IOS XR: 0
- V: Value, set if prefix-SID carries a value (not an index) – IOS XR: always unset
- L: Local, set if prefix-SID has local significance – IOS XR: always unset

Verifying IS-IS Database

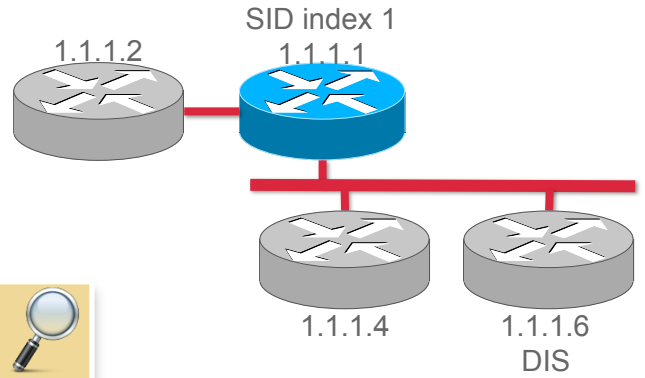


```
<...>

  Metric: 10          IS-Extended xrvr-2.00
    Interface IP Address: 99.1.2.1
    Neighbor IP Address: 99.1.2.2
    ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24002
  Metric: 10          IS-Extended xrvr-6.01
    Interface IP Address: 99.1.6.1
    Neighbor IP Address: 99.1.6.6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24009 System ID:xrvr-6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24007 System ID:xrvr-4
  Metric: 10          MT (IPv6 Unicast) IS-Extended xrvr-2.00
    ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24003
  Metric: 10          MT (IPv6 Unicast) IS-Extended xrvr-6.01
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24001 System ID:xrvr-6
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24011 System ID:xrvr-4
```

IPv4 IS reachability TLV

IPv6 IS reachability TLV

Verifying IS-IS Database

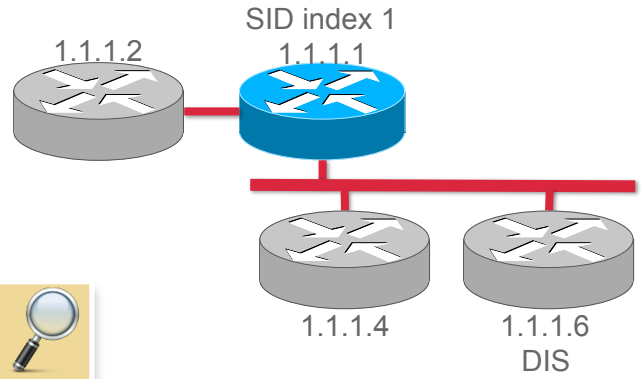


```
<...>

  Metric: 10          IS-Extended xrvr-2.00
    Interface IP Address: 99.1.2.1
    Neighbor IP Address: 99.1.2.2
    ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24002
  Metric: 10          IS-Extended xrvr-6.01
    Interface IP Address: 99.1.6.1
    Neighbor IP Address: 99.1.6.6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24009 System ID:xrvr-6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24007 Syst
  Metric: 10          MT (IPv6 Unicast) IS-Extended xrvr-2.00
    ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24003
  Metric: 10          MT (IPv6 Unicast) IS-Extended xrvr-6.01
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24001 Sys
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24011 Sys
```

IPv4 adjacency SID
F:0, Address-family IPv4
B:0, unprotected
V:1, label value
L:1, local segment
S:0, not a set of adjacencies

IPv6 adjacency SID
F:1, Address-family IPv6
B:0, unprotected
V:1, label value
L:1, local segment
S:0, not a set of adjacencies

Adjacency-SID Flags (see ”Adjacency Segment” slides):

- F: Address-Family, unset: IPv4, set: IPv6
- B: Backup, set if the Adj-SID refers to a protected adjacency (e.g. Using TI-LFA)
- V: Value, set if Adj-SID carries a value – IOS XR: always set
- L: Local, set if Adj-SID has local significance – IOS XR: always set
- S: Set, set if Adj-SID refers to a set of adjacencies – IOS XR: always unset

Verifying IS-IS Database

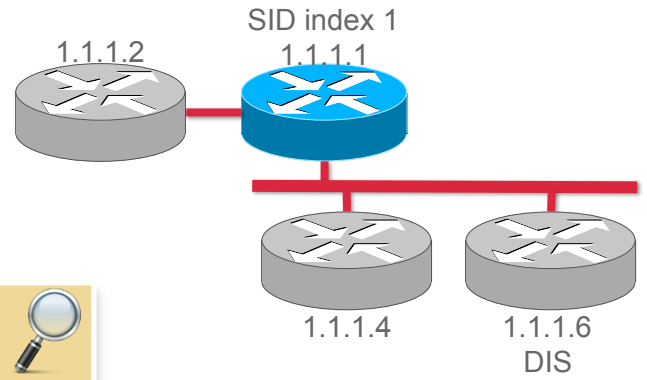


```
<...>

  Metric: 10          IS-Extended xrvr-2.00
    Interface IP Address: 99.1.2.1
    Neighbor IP Address: 99.1.2.2
    ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24002
  Metric: 10          IS-Extended xrvr-6.01
    Interface IP Address: 99.1.6.1
    Neighbor IP Address: 99.1.6.6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24009 System ID:xrvr-6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24007 System ID:xrvr-4
  Metric: 10          MT (IPv6 Unicast) IS-Extended xrvr-2.00
    ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24003
  Metric: 10          MT (IPv6 Unicast) IS-Extended xrvr-6.01
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24001 System ID:xrvr-6
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24011 System ID:xrvr-4
```

IPv4 IS reachability to pseudonode TLV

IPv6 IS reachability to pseudonode TLV

Verifying IS-IS Database

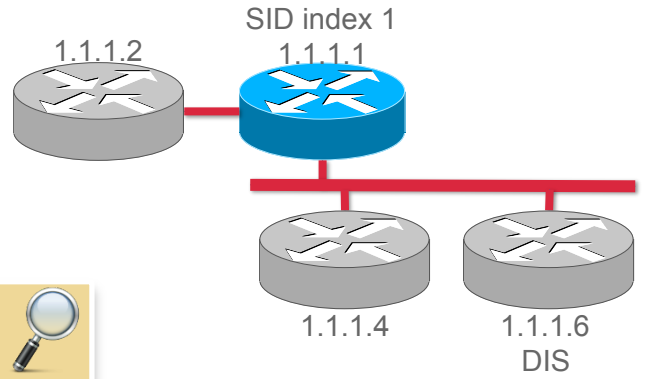


```
<...>

  Metric: 10          IS-Extended xrvr-2.00
    Interface IP Address: 99.1.2.1
    Neighbor IP Address: 99.1.2.2
    ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24002
  Metric: 10          IS-Extended xrvr-6.01
    Interface IP Address: 99.1.6.1
    Neighbor IP Address: 99.1.6.6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24009 System ID:xrvr-6
    LAN-ADJ-SID: F:0 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24007 System ID:xrvr-4
  Metric: 10          MT (IPv6 Unicast) IS-Extended xrvr-2.00
    ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24003
  Metric: 10          MT (IPv6 Unicast) IS-Extended xrvr-6.01
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24001 System ID:xrvr-6
    LAN-ADJ-SID: F:1 B:0 V:1 L:1 S:0 weight:0 Adjacency-sid:24011 System ID:xrvr-4
```

IPv4 LAN adjacency SIDs
F:0, Address-family IPv4
B:0, unprotected
V:1, label value
L:1, local segment
S:0, not a set of adjacencies

IPv6 LAN adjacency SIDs
F:1, Address-family IPv6
B:0, unprotected
V:1, label value
L:1, local segment
S:0, not a set of adjacencies

(LAN) Adjacency-SID Flags (see ”Adjacency Segment” slides):

- F: Address-Family, unset: IPv4, set: IPv6
- B: Backup, set if the Adj-SID refers to a protected adjacency (e.g. Using TI-LFA)
- V: Value, set if Adj-SID carries a value – IOS XR: always set
- L: Local, set if Adj-SID has local significance – IOS XR: always set
- S: Set, set if Adj-SID refers to a set of adjacencies – IOS XR: always unset

OSPF Segment Routing Advertisements

OSPF Configuration Example

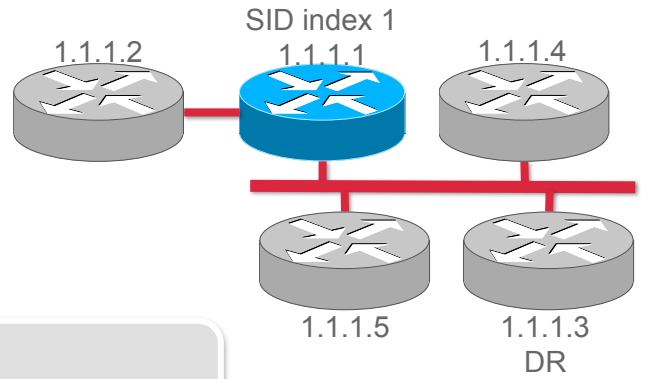


```
router ospf 1
 router-id 1.1.1.1
 segment-routing mpls
 segment-routing forwarding mpls
 area 0
  interface Loopback0
   passive enable
   prefix-sid absolute 16001
  !
  interface GigabitEthernet0/0/0/0
   network point-to-point
  !
  interface GigabitEthernet0/0/0/1
  !
 !
!
```

Enable SR on all areas

Enable SR forwarding on all interfaces

Prefix-SID for loopback0

Adjacency-SIDs will automatically be allocated for adjacencies with SR forwarding enabled

Verifying OSPF Database

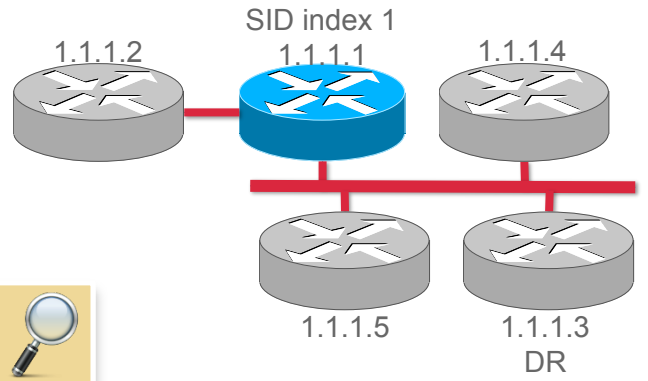


```
RP/0/0/CPU0:xrvr-1#show ospf database self-originate originate


            OSPF Router with ID (1.1.1.1) (Process ID 1)

                Router Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Link count
1.1.1.1         1.1.1.1         13          0x8000000b 0x00b7ba 4

                Type-10 Opaque Link Area Link States (Area 0)

Link ID         ADV Router      Age         Seq#       Checksum Opaque ID
4.0.0.0         1.1.1.1         1085        0x80000003 0x000a66
7.0.0.1         1.1.1.1         1085        0x80000003 0x00e1b2         1
8.0.0.4         1.1.1.1         823         0x80000003 0x0042c0         4
8.0.0.7         1.1.1.1         13          0x80000009 0x0030f4         7
```

Area-scoped (Type10) Opaque LSAs

Router Information Opaque LSA

Opaque LSAs for Segment Routing

Verifying OSPF Database

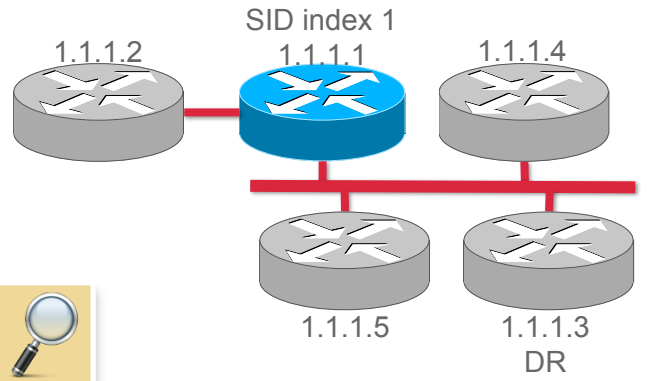


```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 4.0.0.0 self-originate

<...>
  Opaque Type: 4
<...>
     Router Information TLV: Length: 4
     Capabilities:
       Graceful Restart Helper Capable
       Stub Router Capable
       All capability bits: 0x60000000

     Segment Routing Algorithm TLV: Length: 1
       Algorithm: 0

     Segment Routing Range TLV: Length: 12
       Range Size: 8000

         SID sub-TLV: Length 3
          Label: 16000
```

Opaque type 4 (Router Info)

SR algorithm TLV
indicates used algorithms

SRGB TLV

Verifying OSPF Database

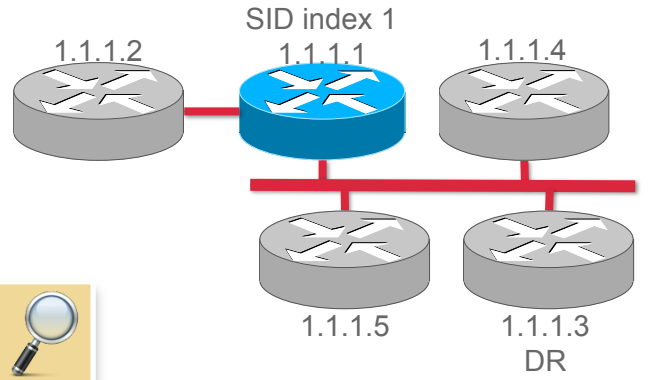


```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 7.0.0.1 self-originate

<...>
 Opaque Type: 7
<...>

    Extended Prefix TLV: Length: 20
      Route-type: 1
      AF        : 0
      Flags     : 0x40
      Prefix    : 1.1.1.1/32

      SID sub-TLV: Length: 8
        Flags     : 0x0
        MTID      : 0
        Algo      : 0
        SID Index : 1
```

Opaque type 7 (Extended Prefix)

N:0, Node-SID

Prefix: intra-area 1.1.1.1/32

Prefix-SID: index 1

Extended Prefix TLV flags (see "Prefix Segment" slides):

- Bit0: A: Attach, set if ABR generates an Extended Prefix TLV for inter-area prefix that is locally connected or attached in another connected area – default in IOS XR: 0
- Bit1: N: Node, set if the prefix identifies the node – default in IOS XR: 1

Prefix-SID sub-TLV flags (see "Prefix Segment" slides):

- Bit1: NP: no-PHP, set if the penultimate hop must NOT pop the prefix-SID before forwarding the packet – default in IOS XR: 0
- Bit2: M: Mapping Server, set if the SID is advertised from the Mapping Server functionality – default in IOS XR: 0
- Bit3: E: Explicit-Null, set if penultimate hop must replace prefix-SID with Explicit-Null label – default in IOS XR: 0
- Bit4: V: Value, set if prefix-SID carries a value (not an index) – IOS XR: always unset
- Bit5:L: Local, set if prefix-SID has local significance – IOS XR: always unset

Verifying OSPF Database

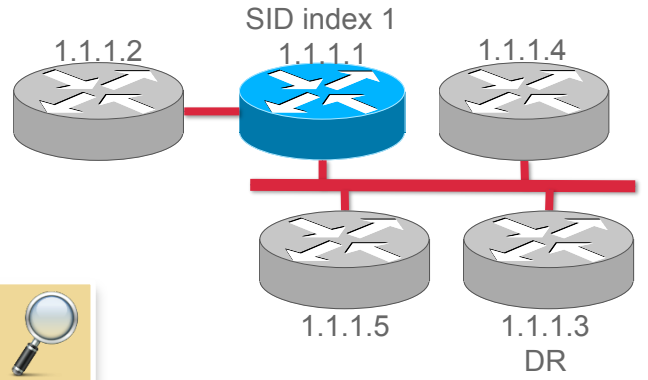


```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 8.0.0.4 self-originate

<...>
 Opaque Type: 8
<...>

 Extended Link TLV: Length: 24
      Link-type : 1
      Link ID   : 1.1.1.2
      Link Data : 99.1.2.1

     Adj sub-TLV: Length: 7
        Flags     : 0x60
        MTID      : 0
        Weight    : 0
        Label     : 24001
```

Opaque type 8 (Extended Link)

Adjacency: p2p to 1.1.1.2

V:1, label value
L:1, local segment

Adjacency-SID: 24001

(LAN) Adjacency-SID sub-TLV flags (see ”Adjacency Segment” slides):

- B: Backup, set if the Adj-SID refers to a protected adjacency (e.g. Using TI-LFA)
- V: Value, set if Adj-SID carries a value (not an index) – IOS XR: always set
- L: Local, set if Adj-SID has local significance – IOS XR: always set
- S: Set, set if Adj-SID refers to a set of adjacencies – IOS XR: always unset

Verifying OSPF Database

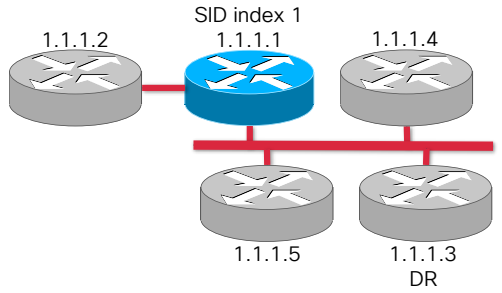


```
RP/0/0/CPU0:xrvr-1#show ospf database opaque-area 8.0.0.7 self-originate
<...>
 Opaque Type: 8
<...>
 Extended Link TLV: Length: 40
      Link-type : 2
      Link ID   : 66.0.0.3
      Link Data : 66.0.0.1

     LAN Adj sub-TLV: Length: 11
       Flags       : 0x60
       MTID        : 0
       Weight      : 0
       Neighbor ID: 1.1.1.4
       Label       : 24002

     LAN Adj sub-TLV: Length: 11
       Flags       : 0x60
       MTID        : 0
       Weight      : 0
       Neighbor ID: 1.1.1.5
       Label       : 24009

     Adj sub-TLV: Length: 7
        Flags      : 0x60
        MTID       : 0
        Weight     : 0
        Label      : 24000
```

Opaque type 8 (Extended Link)

Adjacency: link to transit network

Adjacency-SID for neighbor 1.1.1.4: label 24002

Adjacency-SID for neighbor 1.1.1.5: label 24009

Adjacency-SID to Designated Router: label 24000

Visit us:

cisco.com

segment-routing.net

Acknowledgements:

Ahmed Bashandy
Robert Hanzl
Jose Liste
Steven Luong
Stefano Previdi
Peter Psenak

Thank you.