



Introduction to Segment Routing

Segment Routing (SR) is a flexible, scalable way of doing source routing.

- [Overview of Segment Routing, page 1](#)
- [How Segment Routing Works, page 2](#)
- [Examples for Segment Routing, page 3](#)
- [Benefits of Segment Routing, page 4](#)
- [Segment Routing Global Block, page 6](#)
- [Additional References for Segment Routing, page 8](#)

Overview of Segment Routing

Segment Routing (SR) is a flexible, scalable way of doing source routing. The source chooses a path and encodes it in the packet header as an ordered list of segments. Segments are identifier for any type of instruction. Each segment is identified by the segment ID (SID) consisting of a flat unsigned 32-bit integer. Segment instruction can be:

- Go to node N using the shortest path
- Go to node N over the shortest path to node M and then follow links Layer 1, Layer 2, and Layer 3
- Apply service S

With segment routing, the network no longer needs to maintain a per-application and per-flow state. Instead, it obeys the forwarding instructions provided in the packet.

Segment Routing relies on a small number of extensions to Cisco Intermediate System-to-Intermediate System (IS-IS) and Open Shortest Path First (OSPF) protocols. It can operate with an MPLS (Multiprotocol Label Switching) or an IPv6 data plane, and it integrates with the rich multi service capabilities of MPLS, including Layer 3 VPN (L3VPN), Virtual Private Wire Service (VPWS), Virtual Private LAN Service (VPLS), and Ethernet VPN (EVPN).

Segment routing can be directly applied to the Multiprotocol Label Switching (MPLS) architecture with no change in the forwarding plane. Segment routing utilizes the network bandwidth more effectively than traditional MPLS networks and offers lower latency. A segment is encoded as an MPLS label. An ordered

list of segments is encoded as a stack of labels. The segment to process is on the top of the stack. The related label is popped from the stack, after the completion of a segment.

Segment routing can be applied to the IPv6 architecture with a new type of routing extension header. A segment is encoded as an IPv6 address. An ordered list of segments is encoded as an ordered list of IPv6 addresses in the routing extension header. The segment to process is indicated by a pointer in the routing extension header. The pointer is incremented, after the completion of a segment.

Segment Routing provides automatic traffic protection without any topological restrictions. The network protects traffic against link and node failures without requiring additional signaling in the network. Existing IP fast re-route (FRR) technology, in combination with the explicit routing capabilities in Segment Routing guarantees full protection coverage with optimum backup paths. Traffic protection does not impose any additional signaling requirements.

How Segment Routing Works

A router in a Segment Routing network is capable of selecting any path to forward traffic, whether it is explicit or Interior Gateway Protocol (IGP) shortest path. Segments represent subpaths that a router can combine to form a complete route to a network destination. Each segment has an identifier (Segment Identifier) that is distributed throughout the network using new IGP extensions. The extensions are equally applicable to IPv4 and IPv6 control planes. Unlike the case for traditional MPLS networks, routers in a Segment Router network do not require Label Distribution Protocol (LDP) and Resource Reservation Protocol - Traffic Engineering (RSVP-TE) to allocate or signal their segment identifiers and program their forwarding information.

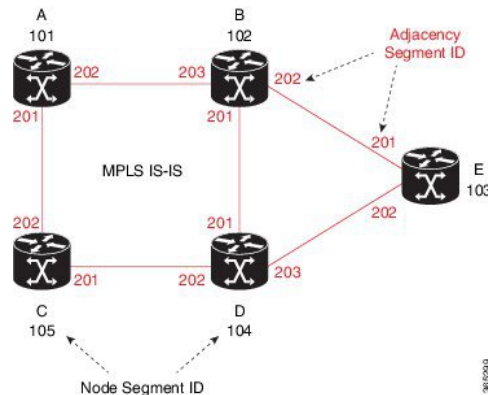
Each router (node) and each link (adjacency) has an associated segment identifier (SID). Node segment identifiers are globally unique and represent the shortest path to a router as determined by the IGP. The network administrator allocates a node ID to each router from a reserved block. On the other hand, an adjacency segment ID is locally significant and represents a specific adjacency, such as egress interface, to a neighboring router. Routers automatically generate adjacency identifiers outside of the reserved block of node IDs. In an MPLS network, a segment identifier is encoded as an MPLS label stack entry. Segment IDs direct the data along a specified path. There are two kinds of segment IDs:

- **Prefix SID**— A segment ID that contains an IP address prefix calculated by an IGP in the service provider core network. Prefix SIDs are globally unique. A node SID is a special form of prefix SID that contains the loopback address of the node as the prefix. It is advertised as an index into the node specific SR Global Block or SRGB.
- **Adjacency SID**— A segment ID that contains an advertising router's adjacency to a neighbor. An adjacency SID is a link between two routers. Since the adjacency SID is relative to a specific router, it is locally unique.

Examples for Segment Routing


The following figure illustrates an MPLS network with five routers using Segment Routing, IS-IS, a label range of 100 to 199 for node IDs, and 200 and higher for adjacency IDs. IS-IS would distribute IP prefix reachability alongside segment ID (the MPLS label) across the network.

Figure 1: An MPLS Network with Five Routers Using Segment Routing



In the previous example, any router sending traffic to router E would push label 103 (router E node segment identifier) to forward traffic using the IS-IS shortest path. The MPLS label-swapping operation at each hop preserves label 103 until the packet arrives at E (Figure 2). On the other hand, adjacency segments behave differently. For example, if a packet arrives at Router D with a top-of-stack MPLS label of 203 (D-to-E adjacency segment identifier), Router D would pop the label and forward the traffic to Router E.

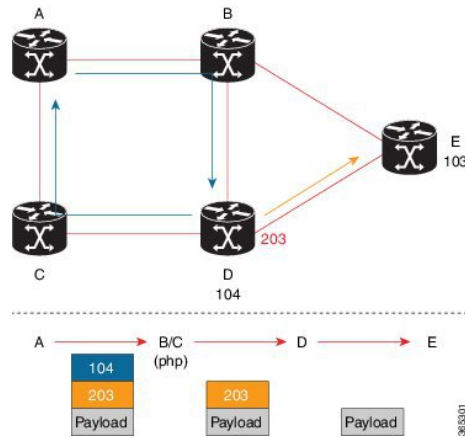
Figure 2: MPLS Label-Swapping Operation



Segment identifiers can be combined as an ordered list to perform traffic engineering. A segment list can contain several adjacency segments, several node segments, or a combination of both depending on the forwarding requirements. In the previous example, Router A could alternatively push label stack (104, 203) to reach Router E using the shortest path and all applicable ECMPs to Router D and then through an explicit interface onto the destination (Figure 3). Router A does not need to signal the new path, and the state information

remains constant in the network. Router A ultimately enforces a forwarding policy that determines which flows destined to router E are switched through a particular path.

Figure 3: Router E Destination Path




Benefits of Segment Routing

- Ready for SDN**—Segment Routing is a compelling architecture conceived to embrace Software-Defined Network (SDN) and is the foundation for Application Engineered Routing (AER). It strikes a balance between network-based distributed intelligence, such as automatic link and node protection, and controller-based centralized intelligence, such as traffic optimization. It can provide strict network performance guarantees, efficient use of network resources, and very high scalability for application-based transactions. The network uses minimal state information to meet these requirements. Segment routing can be easily integrated with a controller-based SDN architecture. Below figure illustrates a sample SDN scenario where the controller performs centralized optimization, including bandwidth admission control. In this scenario, the controller has a complete picture of the network topology and flows. A router can request a path to a destination with certain characteristics, for example, delay, bandwidth, diversity. The controller computes an optimal path and returns the corresponding segment list, such as an MPLS label

stack, to the requesting router. At that point, the router can inject traffic with the segment list without any additional signaling in the network.

Figure 4: SDN Controller



- In addition, segment lists allow complete network virtualization without adding any application state to the network. The state is encoded in the packet as a list of segments. Because the network only maintains segment state, it can support a large number - and a higher frequency - of transaction-based application requests without creating any burden on the network.
- **Simplified**—
 - When applied to the MPLS data plane, Segment Routing offers the ability to tunnel MPLS services (VPN, VPLS, and VPWS) from an ingress provider edge to an egress provider edge without any other protocol than an IGP (ISIS or OSPF).
 - Simpler operation without separate protocols for label distribution (for example, no LDP or RSVP).
 - No complex LDP or IGP synchronization to troubleshoot.
 - Better utilization of installed infrastructure, for lower capital expenditures (CapEx), with ECMP-aware shortest path forwarding (using node segment IDs).
- **Supports Fast Reroute (FRR)**— Deliver automated FRR for any topology. In case of link or node failures in a network, MPLS uses the FRR mechanism for convergence. With segment routing, the convergence time is sub-50-msec.
- **Large-scale Data Center-**
 - Segment Routing simplifies MPLS-enabled data center designs using Border Gateway Protocol (BGP) RFC 3107 - IPv4 labeled unicast among Top-of-the-Rack/Leaf/Spine switches.
 - BGP distributes the node segment ID, equivalent to IGP node SID.
 - Any node within the topology allocates the same BGP segment for the same switch.
 - The same benefits are provided as for IGP node SID: ECMP and automated FRR (BGP PIC(Prefix Independent Convergence)).
 - This is a building block for traffic engineering - SR TE data center fabric optimization.
- **Scalable**—

- Avoid thousands of labels in LDP database.
 - Avoid thousands of MPLS Traffic Engineering LSP's in the network.
 - Avoid thousands of tunnels to configure.
- **Dual-plane Networks—**
 - Segment Routing provides a simple solution for disjointness enforcement within a so-called “dual-plane” network, where the route to an edge destination from a given plane stays within the plane unless the plane is partitioned.
 - An additional SID “anycast” segment ID allows the expression of macro policies such as: "Flow 1 injected in node A toward node Z must go via plane 1" and "Flow 2 injected in node A towards node Z must go via plane 2."
- **Centralized Traffic Engineering—**
 - Controllers and orchestration platforms can interact with Segment Routing traffic engineering for centralized optimization, such as WAN optimization.
 - Network changes such as congestion can trigger an application to optimize (recompute) the placement of segment routing traffic engineering tunnels.
 - Segment Routing tunnels are dynamically programmed onto the network from an orchestrator using southbound protocols like PCE.
 - Agile network programming is possible since Segment Routing tunnels do not require signaling and per-flow state at midpoints and tail end routers.
- **Egress Peering Traffic Engineering (EPE)—**
 - Segment Routing allows centralized EPE.
 - A controller instructs an ingress provider edge and content source to use a specific egress provider edge and specific external interface to reach a destination.
 - BGP “peering” segment IDs are used to express source-routed inter-domain paths.
 - Controllers learn BGP peering SIDs and the external topology of the egress border router through BGP Link Status (BGP-LS) EPE routes.
 - Controllers program ingress points with a desired path.
- **Plug-and-Play deployment—** Segment routing tunnels are interoperable with existing MPLS control and data planes and can be implemented in an existing deployment.

Segment Routing Global Block

Segment Routing Global Block (SRGB) is the range of labels reserved for segment routing. SRGB is local property of an segment routing node. In MPLS architecture, SRGB is the set of local labels reserved for global segments. In segment routing, each node can be configured with a different SRGB value and hence the absolute SID value associated to an IGP Prefix Segment can change from node to node.

The SRGB default value is 16000 to 23999. The SRGB can be configured as follows:

```
Device(config)# router isis 1
Device(config-isis)#segment-routing global-block 45000 55000
```

The SRGB label value is calculated as follows:

- If the platform supports 1000000 labels or more, the SRGB value is from 900000 to $900000 + 2^{16}$.
- If the platform supports less than 1000000 labels, the SRGB value is the last 2^{16} labels.

Restrictions:

- The SRGB size cannot be more than 2^{16} .
- The SRGB upper bound cannot exceed the platform capability.
- The SRGB cannot be configured to be the same value as the default SRGB. So SRGB cannot be configured for 16000 to 23999.

Segment Routing Global Block

This chapter explains the concept of creating a block of labels reserved for a router using segment routing. This block of reserved labels is known as the Segment Routing Global Block (SRGB).

Adjacency Segment Identifiers

The Adjacency Segment Identifier (adj-SID) is a local label that points to a specific interface and a next hop out of that interface. No specific configuration is required to enable adj-SIDs. Once segment routing is enabled over IS-IS for an address-family, for any interface that IS-IS runs over, the address-family automatically allocates an adj-SID towards every neighbor out of that interface.



Note

Only IPV4 address-family supports allocating adj-SIDs.

Prefix Segment Identifiers

A prefix segment identifier (SID) identifies a segment routing tunnel leading to the destination represented by a prefix. The maximum prefix SID value is $2^{16} - 1$.

A prefix SID is allocated from the Segment Routing Global Block (SRGB). The prefix SID value translates to a local MPLS label, whose value is calculated as below:

- If the platform supports 1000000 labels or more, then the MPLS label corresponding to the prefix SID value is $900000 + sid\text{-value}$.
- If the platform supports less than 1000000 labels, then the MPLS label corresponding to the prefix SID value is $maximum\text{-supported-label-value} - 2^{16} + sid\text{-value}$.

When a prefix SID value x is configured, the prefix SID translates to a label value equivalent to $x +$ lower boundary of SRGB. For example, in the platform supporting 1000000 MPLS labels or more if the default

SRGB is used, configuring a prefix-SID of 10 for interface Loopback 0 with IPv4 address 1.0.0.1/32 results in assigning the label 9000010 16010 to the prefix 1.0.0.1/32.

BGP Prefix Segment Identifiers

Segments associated with a BGP prefix are known as BGP Prefix-SIDs.

- BGP Prefix-SIDs are always global within a Segment Routing or BGP domain
- BGP Prefix-SIDs identifies an instruction to forward the packet over ECMP-aware best path computed by BGP for a given prefix

Segment Routing requires BGP speaker to be configured with a Segment Routing Global block (SRGB). Generally, SRGB is configured as a range of labels, SRGB = [SR_S, SR_E].

- SR_S = Start of the range
- SR_E = End of the range

Each prefix is assigned with its own unique label index.

In the following example, a BGP route policy, set label index, is defined using the route-policy **name** command.

Configure the Segment Routing Global Block (SRGB) in BGP. If the route label path has a label-index attribute and SRGB is configured, then local label route is allocated from SRGB. If label-index is added to redistributed routes using route-policy, then BGP presents label-index as an attribute with the route.

```
router bgp 100
  bgp log-neighbor-changes
  neighbor 192.0.2.1 remote-as 100
  neighbor 192.0.2.1 update-source Loopback0
  neighbor 192.0.23.3 remote-as 300
  !
  address-family ipv4
    segment-routing mpls
    neighbor 192.0.2.1 activate
    neighbor 192.0.2.1 send-label
    neighbor 192.0.23.3 activate
  exit-address-family
```

Additional References for Segment Routing

Related Documents

| Related Topic | Document Title |
|---------------|--|
| Videos | <ul style="list-style-type: none"> • Introduction to Cisco Segment Routing (YouTube) • Introduction to Cisco Segment Routing (CCO) |