Introduction to Segment Routing for the MPLS Dataplane

This chapter introduces the concept of Segment Routing and contains the following sections:

Feature Information for Segment Routing

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to https://cfnng.cisco.com/. An account on Cisco.com is not required.

Table 1: Feature Information for Segment Routing

Feature Name	Releases	Feature Information
Introduction to Segment Routing	Cisco IOS XE Amsterdam 17.3.2	Segment Routing is a flexible, scalable way of doing source routing.

Overview of Segment Routing for the MPLS Dataplane

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 an identifier for any type of instruction. Each segment is identified by a segment ID (SID) . A 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 a specific set of links.

• 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 extensions to the Intermediate System-to-Intermediate System (IS-IS) and Open Shortest Path First (OSPF) protocols. When operating using the MPLS (Multiprotocol Label Switching) dataplane 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).

Note   The Cisco IOS XE Amsterdam 17.3.2 release only provides support for the MPLS dataplane for the IPv4 address family.

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 provides automatic traffic protection without any topological restrictions. The network protectstraffic 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 an explicit path, or an Interior Gateway Protocol (IGP) calculated shortest path. Segments represent subpaths that a router can combine to form a complete path to a network destination. Each segment has an identifier (Segment Identifier) that is advertised throughout the network using IGP extensions. 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 support forwarding.

Each router (node) and each link (adjacency) have an associated segment identifier (SID). Node segment identifiers are globally unique. 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 to the advertising node and represents a specific adjacency, such as an egress interface, to a neighboring router. Routers automatically generate adjacency identifiers outside of the reserved block used for node IDs. In an MPLS network, a segment identifier is encoded as an MPLS label stack entry. Segment IDs direct data packets along a specified path. There are two kinds of segment IDS:

• **Prefix SID**:

A segment ID that is associated with an IP address prefix associated with a node in the network. Prefix SIDs are globally unique. A node segment is a special prefix segment that is bound to the loopback address of a node. It is advertised as an index into node specific SR Global Blocks (SRGBs).

• **Adjacency SID**:

A segment ID that is associated with a link between two routers. Adjacency SIDs are scoped to a specific router.

A node segment can be a multihop path while an adjacency segment is always a one-hop path.

Examples for Segment Routing

The following figure illustrates an MPLS network with five routers using Segment Routing, IS-IS, a label range of 16000 to 23999 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 Figure 1, any router sending traffic to router E would push label 16103 (router E node segment identifier) to forward traffic using the IS-IS shortest path. The MPLS label-swapping operation at each hop preserves label 16103 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
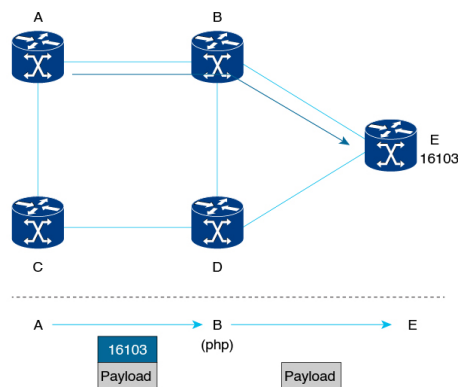
Segment identifiers can be combined as an ordered list to perform traffic engineering. A segment list can contain several adjacency segments, several node segments, or a combination of both depending on the forwarding requirements. In the previous example, Router A could alternatively push label stack (16104, 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
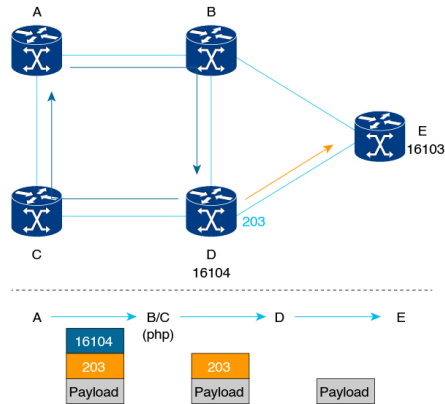

Benefits of Segment Routing

- **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.

  The following 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.

  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.

Figure 4: SDN Controller

- **Simplified Operation**:

  - 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.

- **Dual-Plane Networks**:

  - Segment Routing provides a simple solution for disjointness enforcement within a *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

The Segment Routing Global Block (SRGB) is the range of labels reserved for segment routing globally scoped SIDs. SRGB is local property of a segment routing node. In the 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 range is 16000 to 23999. The SRGB can be configured as follows:

```
Device(config)# segment-routing mpls
Device(config-srmpls)#segment-routing global-block 45000 55000
```

The SRGB label value is calculated hop-by-hop as follows:

- Node-SIDs are advertised as an index into the local SRGB.

- IGPs calculate the appropriate MPLS label associated with a Node SID by adding the index to the SRGB advertised by the next hop.

- Platforms may have specific limitations on the lower and upper bounds for the SRGB. These bounds will be visible in the help string associated with the segment-routing global-block command e.g.:

  ```
  Router(config-srmpls)#global-block ?
    <16-1048575>  SR GB/LB Label Range Start
  ```

Adjacency Segment Identifiers

An Adjacency Segment Identifier (adj-SID) is a local label that directs packets to a specific interface and a next hop. No specific configuration is required to enable adj-SIDs. Adjacency SIDs will be automatically allocated by the IGP for every neighbor.

Prefix Segment Identifiers

A prefix segment identifier (SID) is used to forward traffic to a destination represented by a prefix.

A prefix SID is an index into the Segment Routing Global Block (SRGB). The index maps to a local MPLS label, whose value is calculated for each node by adding the index value to the locally advertised SRGB starting value. For example:

- If a node advertises an SRGB range of 16000 – 23999, a SID index of 100 would map to MPLS label 16100.

- If a node advertises an SRGB range of 25000 - 39999, a SID index of 100 would map to MPLS label 25100

Additional References for Segment Routing

Related Documents

Related Topic	Document Title
Videos	• Introduction to Cisco Segment Routing (YouTube) • Introduction to Cisco Segment Routing (CCO)