White Paper

Segment Routing: Prepare Your Network for New Business Models

Segment Routing: Introduction and Value Proposition

Segment routing is a network technology focused on addressing the pain points of existing IP and Multiprotocol Label Switching (MPLS) networks in terms of simplicity, scale, and ease of operation. It's a foundation for application engineered routing because it prepares the networks for new business models where applications can direct network behavior.

Segment routing seeks the right balance between distributed intelligence and centralized optimization and programming. It was built for the software-defined networking (SDN) era.

Segment routing enables enhanced packet forwarding behavior. It enables a network to transport unicast packets through a specific forwarding path, other than the normal shortest path that a packet usually takes. This capability benefits many use cases, and you can build those specific paths based on application requirements.

Segment routing uses the source routing paradigm. A node, usually a router but it can also be a switch, a trusted server, or a virtual forwarder running on a hypervisor, steers a packet through an ordered list of instructions, called segments. A segment can represent any instruction, topological or service-based. A segment can have a local semantic to a segment-routing node or global within a segment-routing network. Segment routing allows you to enforce a flow through any topological path and service chain while maintaining per-flow state only at the ingress node to the segment-routing network. To be aligned with modern IP networks, segment routing supports equal-cost multipath (ECMP) by design, and the forwarding within a segment-routing network uses all possible paths, when desired.

Segment routing relies on a small number of extensions to routing protocols (Intermediate System-to-Intermediate System [IS-IS], Open Shortest Path First [OSPF], and Border Gateway Protocol [BGP]), and can operate with either an MPLS or an IPv6 data plane. All the currently available MPLS services, such as Layer 3 VPN (L3VPN), L2VPN (Virtual Private Wire Service [VPWS], Virtual Private LAN Services [VPLS], Ethernet VPN [E-VPN], and Provider Backbone Bridging Ethernet VPN [PBB-EVPN]), can run on top of a segment-routing transport network.

One of the key characteristics of segment routing is simplicity:

- From a configuration perspective, the number of lines required to enable segment routing is minimum, usually three lines of configuration.
- From an operational perspective, it simplifies the operation of an MPLS network by making the label value constant across the core of the network.
- From a scale and simplicity perspective, segment routing is especially powerful in the era of SDN with application requirements programming the network behavior and where traffic differentiation and engineering are done at a finer granularity (for example, application-specific).

In existing traffic-engineering solutions, specifically Resource Reservation Protocol - Traffic Engineering (RSVP-TE), each path, when computed, needs to be signaled, and the state for each path must be present in each node traversed by the path. Segment routing allows you to implement traffic engineering without a signaling component.

Therefore, its architecture scales significantly more, and at the same time it simplifies the hardware requirements for the routers in the core of the network (for example, P routers). In addition, if ECMPs are present in the network, segment-routing traffic-engineering tunnels can use all the paths for load balancing flow across them, whereas if RSVP-TE is used to build a traffic-engineering tunnel, only one path is selected.

In such a demanding context of application dictating network behavior, with traffic differentiation done for millions of flows, simplicity and scalability provided by segment routing matter more than ever. Also it's important to note that the use of a segment-routing architecture end to end, from the data center through the WAN, allows smooth, service-rich integration between the data center and the WAN so a communication policy for the applications running in the data center can be implemented end to end.

In a segment-routing network using the MPLS data plane, Label Distribution Protocol (LDP) and RSVP-TE signaling protocols are not required; instead label distribution is performed by the Interior Gateway Protocol (IGP) (IS-IS or OSPF) or BGP. Removing protocols from the network simplifies its operation and makes it more robust and stable by eliminating the need for protocol interaction.

Another benefit that segment routing brings is automated and native Fast Reroute (FRR) capability, with sub-50-millisecond convergence time. FRR has been deployed in order to cope with link or node failures. Segment routing supports FRR on any topology, without any additional signaling protocol, and it supports node and link protection. In a segment-routing network the FRR backup path is optimal because it is provided over the postconvergence path, avoiding transient congestion and suboptimal routing while simplifying the operation and deployment.

Segment routing allows incremental and selective deployment without any requirement of a "flag day" or massive upgrade of all network elements; you can deploy and integrate it with existing MPLS networks because it is fully interoperable with the existing MPLS control and date planes.

Segment routing is currently undergoing standardization at the IETF. A Cisco innovation, it can be seen on the Segment Routing Architecture IETF Draft (https://datatracker.ietf.org/doc/draft-ietf-spring-segment-routing/). It is supported by several industry leaders.

The Cisco® ASR 9000 Series Aggregation Services Router was the first router in the industry to support segment routing; it started supporting this technology with Release 5.2.0, made available in July 2014.

Market Needs and Target Audience for Segment Routing

Market Needs

Traditionally, traffic steering within networks has been based primarily on an IP lookup at each router lying on the path toward the destination. The underlying IGP protocol was used to distribute the topology information and calculate the shortest paths to be followed from the ingress to the egress node. In the last decade, the dependence of many businesses on network performance and availability became more and more relevant. Strict service-level agreements (SLAs) in terms of packet loss, delay, jitter, and available bandwidth became a key business differentiator. Such new requirements pushed network evolution toward architectures that allow for steering traffic with more flexibility.

In the 1990s, the MPLS architecture introduced a performant tunneling mechanism and a traffic steering function, modeled on the ATM/Frame Relay architecture. The tunneling function was fundamental for the success of the MPLS technology, because it enabled the introduction of IP-/MPLS-based VPNs.

The RSVP-TE-based traffic-engineering function of MPLS has not enjoyed the same success and is rarely deployed compared to the VPN features, for three fundamental reasons. First, the poor balancing characteristics of the ATM/Frame Relay model do not fit the true nature of IP, which is based on networks offering abundant ECMPs. To overcome this fundamental incoherence, a notorious number of MPLS RSVP-TE tunnels would need to be replicated, thus elevating the difficulty of managing and monitoring the network. Second, control- and data-plane scalability problems are caused by the state required at each hop along any explicit path. Third, the RSVP-TE deployment model observed until now is based on distributed computation, leading to unpredictable placement of the traffic, nonoptimal use of the resources, and slow reoptimization.

From this experience, network operators are now requiring an evolution toward flexible, operationally friendly, incrementally deployable, yet scalable and simple network architectures. The emergence of application-centric networking and cloud-based services further highlights the need for such an evolution.

Segment routing was created to address the needs described previously, as well as to evolve the network infrastructure to the SDN era.

Target Audience for Segment Routing

Segment routing supports use cases with applicability for service providers, over-the-top (OTT) providers, content and web providers, and large enterprises across the WAN, metropolitan-area network (MAN), and data center.

Segment routing is defined and applicable for unicast traffic. The application of the source-route concept to IP Multicast is not in the scope of segment routing.

This document focuses mostly on the applicability of segment routing to the WAN, with an MPLS data plane. The focus is on the scenarios where the customer operates the end-to-end network, meaning in an MPLS network for which the operator manages the provider-edge router (PE router) as well as the P routers on the network.

Basic Architecture of Segment Routing

Segment routing is a network technology that provides enhanced packet-forwarding behavior while minimizing the need for maintaining awareness of mass volumes of network state. The following is an extract from the IETF draft defining Segment Routing:

“Segment Routing (SR) leverages the source routing paradigm. A node steers a packet through an ordered list of instructions, called segments. A segment can represent any instruction, topological or service-based. A segment can have a local semantic to an SR node or global within an SR domain. Segment routing allows you to enforce a flow through any topological path and service chain while maintaining per-flow state only at the ingress node to the SR domain.

Segment routing can be directly applied to the MPLS architecture with no change on the forwarding plane. 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. Upon completion of a segment, the related label is popped from the stack.

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. Upon completion of a segment, the pointer is incremented.”

The following sections detail the architecture of segment routing.

Concepts of Segment Routing

Segment

A segment is a 32-bit identification for either a topological instruction or a service instruction. A segment can be either global or local. The instruction associated with a global segment is recognized and executed by any segment-routing-capable node in the domain. The instruction associated with a local segment is supported by only the specific node that originates it.

The following instructions could be used in a segment:

- Service
- Context
- Locator
- IGP-based forwarding construct
- BGP-based forwarding construct
- Local value or global index

Global Segment

Any node in the segment-routing domain understands the instruction associated with a global segment. Any node in the domain installs the related instruction in its Forwarding Information Base (FIB). Global segments fall in a subspace of the segment (or label) space called the Segment Routing Global Block (SRGB). The SRGB is usually defined as the range 16000 to 23999, and all the nodes in a network are allocated the same SRGB; this stipulation is important to fulfill the requirement for operational simplification. Note that the use of a common SRGB in all nodes is not a requirement; a different SRGB at every node can be used, if needed.

Local Segment

The instruction associated with a local segment is supported by only the node originating it. No other node installs a remote local segment in its FIB.

For example: If node N allocates segment 29001 to the local forwarding instruction "complete the segment and forward the packet onto interface I", then it advertises this local instruction with absolute value 29001. No other node installs that segment in its segment-routing FIB and hence no conflict can arise.

IGP Segments

IGP (OSPF or IS-IS) can allocate segments for different purposes.

Prefix Segment

This global segment is signaled by IS-IS or OSPF associated to a prefix that is used to steer traffic along the ECMP-aware shortest path to the prefix.

Node Segment

A node segment is a segment allocated to a loopback that identifies a specific node.

In this example (Figure1), 16005 is the node segment of router 5, in other words, it's a segment allocated to a loopback of router 5, so a packet injected anywhere in the network with top segment 16005 will reach 5 by the shortest-path. The imposition of the segment happens on the ingress router, in the Figure 1 at Router 1.

Figure 1. Node Segment Imposition

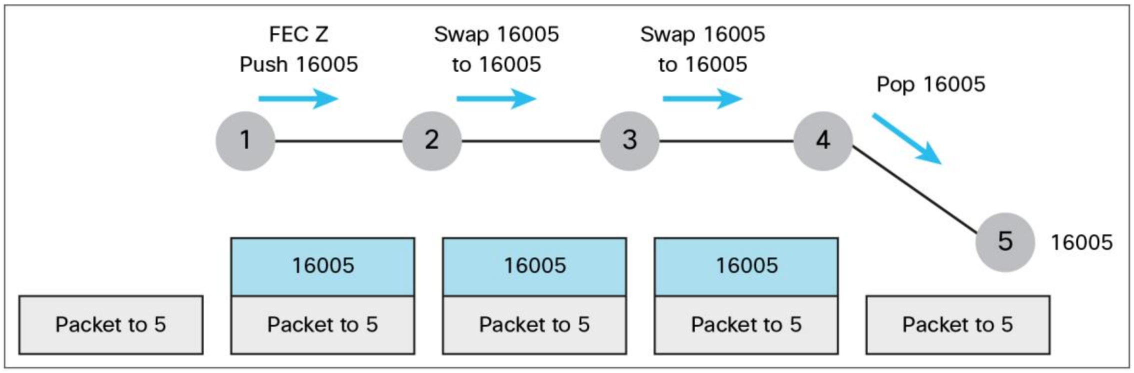

Adjacency Segment

An adjacency segment is a local segment signaled by IS-IS or OSPF; it is used to steer traffic onto an adjacency or a set of adjacencies.

In the example shown in Figure 2, node 5 allocates a local segment 29057 and maps it to the instruction "complete the segment and forward along the interface 57". Node 5 advertises the adjacency segment in IGP. Node 5 is the only node to install the adjacency segment in FIB. A packet received by node 5 with segment 29057 is forced through link 57.

Figure 2. Adjacency Segment

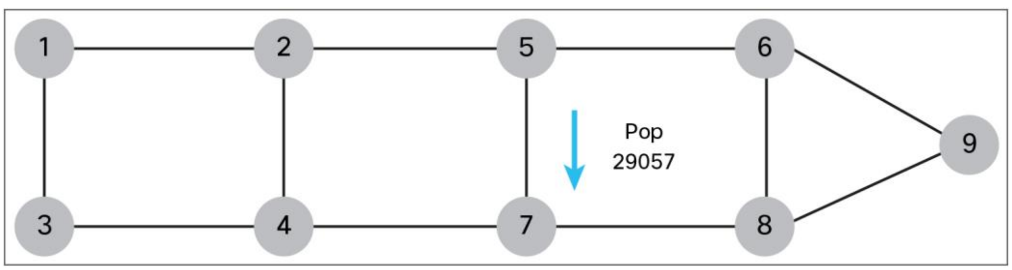

BGP Prefix Segment

The BGP prefix segment is a global segment signaled by BGP associated to a prefix that is used to steer traffic along the ECMP-aware shortest path to the BGP prefix. In Figure 3, node 1 advertises a BGP prefix with segment 16001, and when node 12 needs to send traffic toward node 1, it can use the shortest path to the BGP prefix advertised by node 1.

Figure 3. BGP Prefix Segment

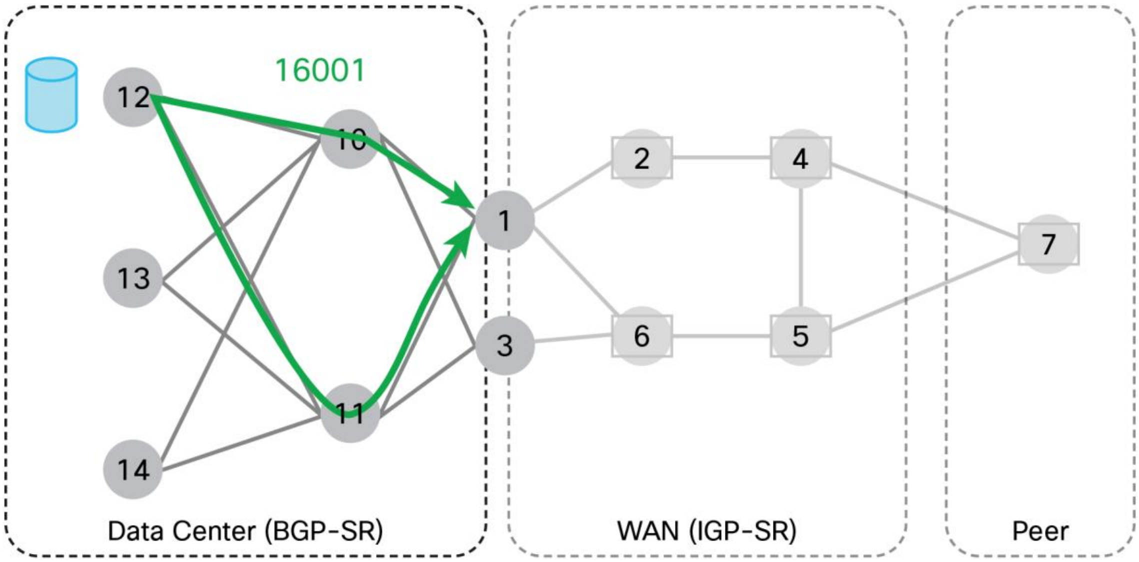

BGP Peering Segment

The BGP segment is a local segment signaled by BGP Link State (topology information) to a SDN controller; it is used to steer traffic onto a BGP peer or over specific links. In the example in Figure 4, to reach node 7 the SDN controller can consider whether or not the application requires a lower-latency path through link 4 to link 7 or higher latency or higher bandwidth, which would be translated into the segment of link 5 to link 7.

Figure 4. BGP Peering Segment

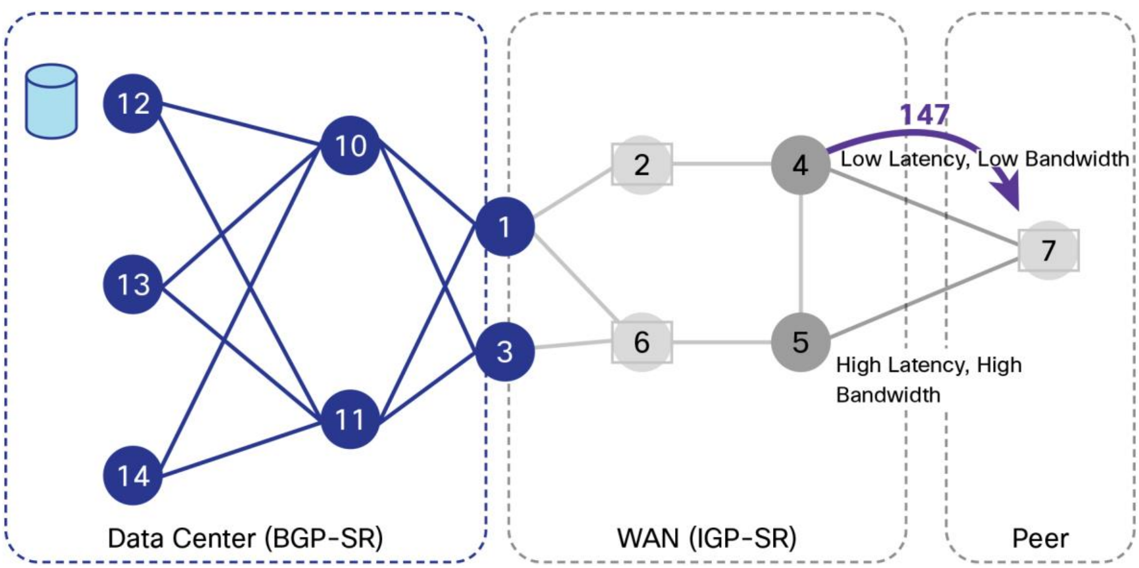

Segment Operation in MPLS Data Plane

You can apply segment routing directly to the MPLS architecture with no change on the forwarding plane. 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. Upon completion of a segment, the related label is popped from the stack.

Segment-routing MPLS data-plane operations are push, swap, and pop as per the traditional MPLS forwarding.

Figure 5 illustrates the MPLS data-plane operations.

Figure 5. MPLS Data-Plane Operations

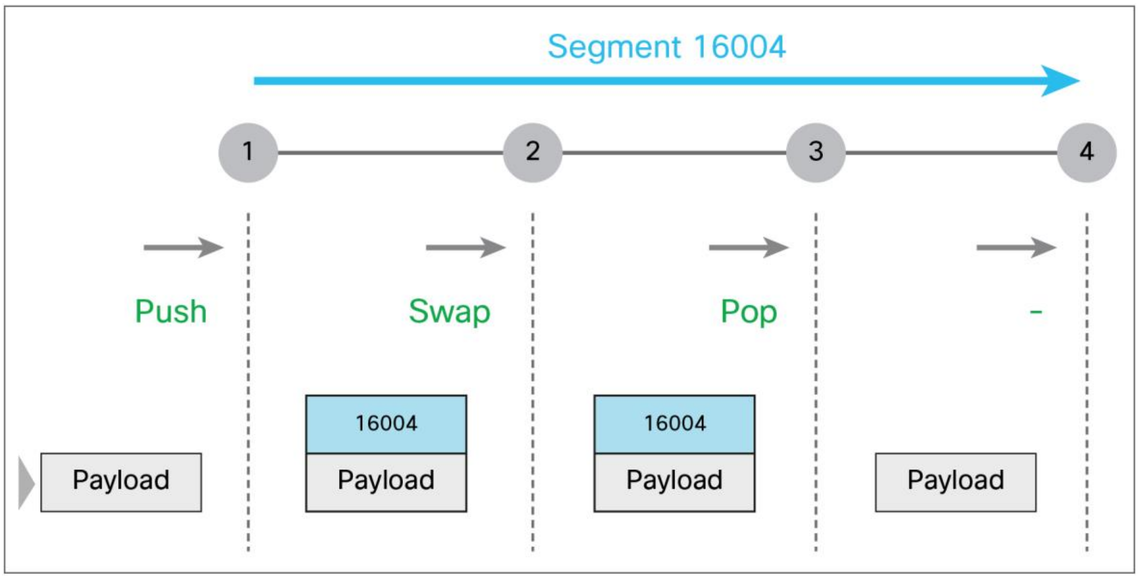

In the figure, Router 4 advertises its loopback prefix 4.4.4/32 with attached prefix session ID 16004. Router 1, the headend pushes label 16004 and program the egress outgoing interface.

Router 2 will have an entry on its LFIB for the remote prefix-SID of value 16004; it swaps operation and program the egress outgoing interface. Router 3 also has an entry on its LFIB for the remote prefix-SID of value 16004, considering penultimate hop behaviour (PHP), Router 3 does a pop operation on the label and send the packet on the programmed egress interface.

Finally router 4 receives the packet without a label and forwards it completed based on the IP address or service label information if one was present.

Segment-Routing IGP-Based Control Plane

Extensions to both IS-IS and OSPF protocols have been made to support segment routing.

Segment-Routing IS-IS Control Plane

Intermediate System to Intermediate System TLV extensions have been implemented for segment-routing support in IS-IS. The implementation is based on the IETF draft “draft-ietf-IS-IS-segment-routing-extensions”.

The following IS-IS TLV extensions have been implemented in Cisco IOS® XR Software:

- Extended IP reachability TLV (135): Prefix-SID sub-TLV
- IS-IS Router Capability TLV (242): Segment routing Capability sub-TLV
- Extended IS Reachability TLV (22): Adjacency-SID sub-TLV
- Mapping Server SID Mappings TLV (149): SID/Label Binding TLV

Segment-Routing IS-IS Basic Configuration

The following is an example of enabling segment routing with IS-IS as the IGP.

As a basic requirement, IS-IS wide metrics must be used and segment routing must be enabled under the IPv4 unicast address family. Then a prefix-sid is assigned for a loopback interface to be the node segment.

```
router ISIS 1
 address-family ipv4 unicast
  metric-style wide
  segment-routing mpls
 !
 interface Loopback0
  address-family ipv4 unicast
   prefix-sid {absolute|index} {<SID value>|<SID index>}
```

Segment-Routing OSPF Control Plane

OSPF extensions have been implemented to support segment routing. The implementation is based on the “draft-psenak-ospf-segment-routing-extensions”.

OSPF Basic Configuration

The following is an example of enabling segment routing with OSPF as the IGP.

As a basic requirement, segment routing for OSPF must be enabled under the OSPF instance or OSPF area, and segment-routing forwarding must be enabled under instance, area(s), or interface(s). The usual command inheritance is applied because the segment-routing commands have interface and area scope.

```
router ospf 1
 segment-routing mpls
 segment-routing forwarding mpls
 area 0
  interface Loopback0
   prefix-sid index 1
```

Overview of Segment-Routing Use Cases

This section describes some of the segment-routing use cases. It’s important to note that segment-routing use cases consistently use the same building blocks; that is, segment types, described in previous subsections of the section “Basic Architecture of Segment Routing”.

The use cases covered here are not an exhaustive list of segment-routing use cases.

Simpler and Efficient Transport of MPLS Services

Segment routing, applied to the MPLS data plane, offers the ability to tunnel MPLS services (L3VPN, VPLS, and VPWS) from an ingress provider edge to an egress provider edge, without any other protocol than IS-IS or OSPF. LDP and RSVP-TE signaling protocols are not required.

The operator needs only to allocate one node segment per provider edge, and the segment-routing IGP control plane automatically builds the required MPLS forwarding constructs from any provider edge to any provider edge.

This use case is equally applicable for IPv4 VPN or IPv6 VPN with transport over IPv4 segments.

Figure 6 shows segment routing enabled on the provider edge and P routers (R1 to R6). R1 and R2 are provider-edge routers and have a Multiprotocol BGP (MP-BGP) session between them where they exchange VPNv4 and VPNv6 routes and labels. The forwarding within the segment-routing domain (R1 to R6) is done based on the segment-routing segment toward R2 or R1, depending on the direction the traffic is flowing and the ECMP-aware shortest path. The figure shows how segment routing is used as the transport label, instead of LDP or RSVP-TE, and the MP-BGP label is still used for the L3VPN service.

Figure 6. Transporting MPLS L3VPN Services over Segment Routing

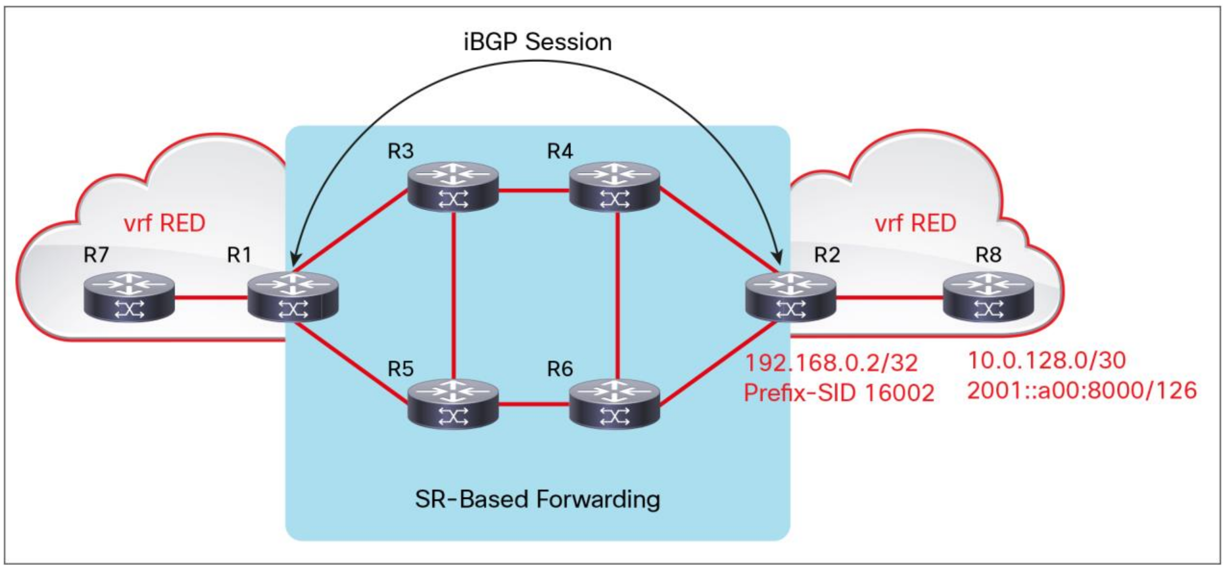

Supporting MPLS services (L3VPN, VPLS, and VPWS) with segment routing has the following benefits:

- Simple operation: You have only one intradomain protocol to operate: IGP. There is no need to support or configure IGP or LDP synchronization extensions.
- Excellent scaling: Only one node SID per provider edge is required.

Fast Reroute: Segment-Routing TI-LFA

Segment routing aims at supporting services with tight SLA guarantees. To meet this goal, the ability to fast reroute after the sudden failure of network components, specifically links or nodes, is mandatory.

Topology-Independent Loop-Free Alternate Fast Reroute (TI-LFA) is used within a segment-routing deployment to achieve this goal; TI-LFA provides 50-msec convergence time. It also provides 100-percent coverage guaranteed in any IGP network.

Fast Reroute (TI-LFA) with segment routing is simple to operate and understand because the protected path is automatically computed by the IGP, without the need of LDP or RSVP-TE. It prevents transient congestion and suboptimal routing by using the postconvergence path, which is the optimum path. The capacity of the postconvergence path is typically planned by network architects to support the postconvergence routing of the traffic for any expected failure, so if this most desired path (postconvergence path) is used, there is much less need for the operator to tune the decision among which protection path to choose. With segment-routing TI-LFA the protection path will automatically follow the natural backup path that would be used after local convergence, helping reduce the amount of path changes and hence service transients: one transition (preconvergence to postconvergence) instead of two (preconvergence to FRR and then postconvergence). With segment-routing TI-LFA there is no midpoint backup state; the only device that keeps state is one that suffered a failure, and it encodes the backup path through a list of segments.

It's important to note that TI-LFA protects IP, segment routing, and LDP-labeled traffic, and you can deploy it incrementally.

Segment Routing and LDP Coexistence and Smooth Interworking

The MPLS architecture permits concurrent usage of multiple label distribution protocols such as LDP, RSVP-TE, and segment-routing control plane. Segment routing uses a dedicated label space (SRGB): all the segment-routing global labels are allocated from this block. All dynamic labels are uniquely allocated outside the SRGB block, allowing segment routing and LDP to coexist without interaction ("ships in the night"). If both LDP and segment routing are enabled on the same node, LDP is preferred by default; the preference segment routing or LDP is configurable through the command-line interface (CLI).

Segment routing is fully interoperable with the LDP and MPLS data planes, providing investment protection and allowing incremental adoption. The mapping-server function advertises the Prefix-to-SID mappings on behalf of non-segment-routing-aware nodes.

The following four deployment models are fully supported (Figure 7):

- Segment routing to LDP: Use this deployment model when a node is segment routing-capable but its next hop along the shortest path to the destination is not. In this case, the prefix segment is connected to the LDP label switched path.
- LDP to segment routing: Use this deployment model when a node is LDP-capable but its next hop along the shortest path to the destination is not. In this case, the LDP LSP is connected to the prefix segment; this connection is done automatically.
- Segment routing over LDP (segment routing to LDP followed by LDP to segment routing): At the segment routing/LDP boundary, the segment-routing prefix segment is mapped to an LDP LSP. At the LDP/segment routing boundary, the LDP LSP is mapped to a segment-routing prefix segment.
- LDP over segment routing (LDP to segment routing followed by segment routing to LDP): At the LDP/segment-routing boundary, the LDP LSP is mapped to a segment-routing prefix segment. At the segment-routing/LDP boundary, the segment-routing prefix segment is mapped to an LDP LSP.

Figure 7. Deployment Models

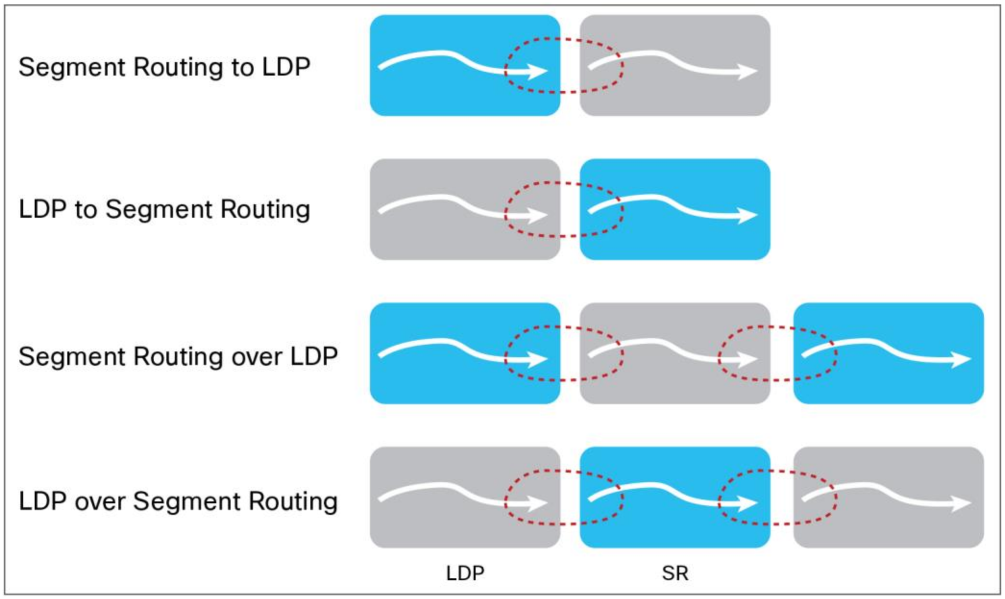

Traffic Engineering

Segment-routing traffic-engineering use cases will vary based on the customer's technical and business needs; for example, some customers have regulatory or business requirements to provide disjoint paths within their network, whereas others need to steer traffic with heavy load away from the shortest path toward a higher-latency, higher-bandwidth path.

Segment routing enables you to steer the traffic in any arbitrary manner within the network, so you can steer traffic along ECMP-aware paths or express deterministic non-ECMP paths.

This section describes some of the segment-routing traffic-engineering use cases. The use of segment-routing traffic engineering is not limited to the use cases covered herein, of course, and in all of the segment-routing traffic-engineering use cases the following properties are maintained:

- There is no midpoint state (n^2 scale in RSVP-TE)
- There is no extra protocol (RSVP-TE)
- There is support for Native ECMP
- Few segments are required
- Segment-routing traffic engineering supports distributed or centralized computation

Disjoint Paths in the WAN: Simply Achievable with Segment-Routing Anycast Segment

The segment-routing architecture defines an anycast segment as a segment attached to an anycast IP prefix. The anycast node segment is an interesting tool for traffic engineering because it makes it easy to express macro traffic-engineering policies, such as "go via plane1 of a dual-plane network" or "go via Region Europe".

Many WAN networks are built according to a dual-plane design; Figure 8 is a representation of a dual-plane network, where we have plane blue and plane red.

We assume a common design rule found in such deployments: the interplane link cost/metric (for example, router 41 to router 31) is set such that the route to an edge destination from a given plane stays within the plane unless the plane is partitioned.

On the network shown in Figure 8, the four blue routers (31, 32, 33 and 34) are configured with an anycast loopback address 192.0.2.1/32 and an Anycast-SID 16111.

Figure 8. Dual-Plane Network

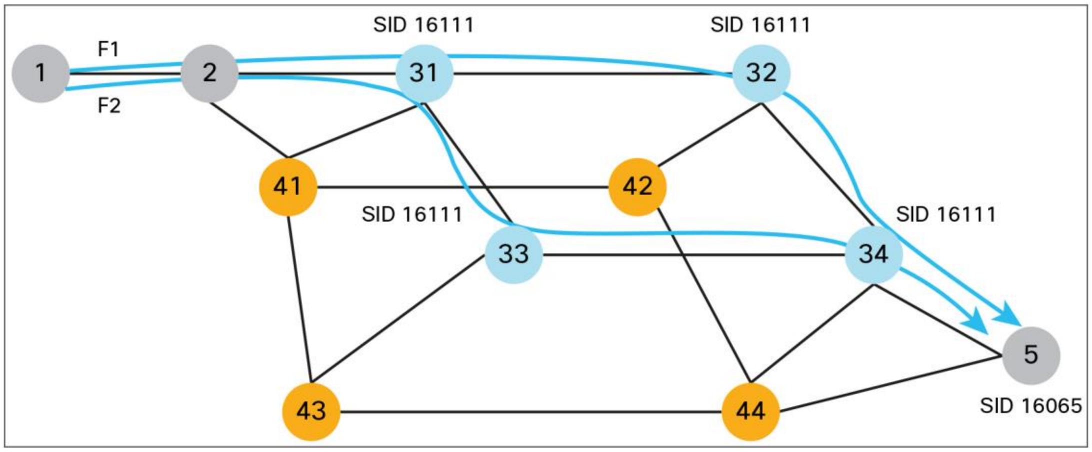

Let us assume a bank or a government customer has requested that certain flows F1, F2, F3, and F4 injected at router 1 and destined to router 5 should be transported across disjoint paths.

Segment routing makes it extremely easy and simple to meet this requirement; you need only to make a configuration on router 1 to direct F1 and F2 packets (represented on the figure) toward a traffic-engineering tunnel that pushes segments 16111, 16065 into the packets. When router 2 receives the packet, it forwards it to plane blue because segment 16111 is the top segment; after the packet enters the blue plane, it follows the ECMP-aware shortest path to router 5 (segment 16065).

Router 1 would be configured with a second tunnel that pushes the anycast segment for plane red (not represented on the figure).

This simple example shows how you can use segment-routing traffic engineering to provide disjoint paths; it's worth noting that the only device in the network on which the policy is configured is router 1; no other device, including router 2, has any state or configuration to fulfill this use case.

Path Avoidance

Segment routing supports distributed or centralized computation. On the example of path avoidance shown in Figure 9, a centralized SDN controller dynamically computes the path to steer the traffic through the lowest-cost path and to avoid the blue link, the SDN controller programs router 1 with the associated segment list. Router 1 is the only node in the network to maintain state.

Figure 9. Example of Path Avoidance

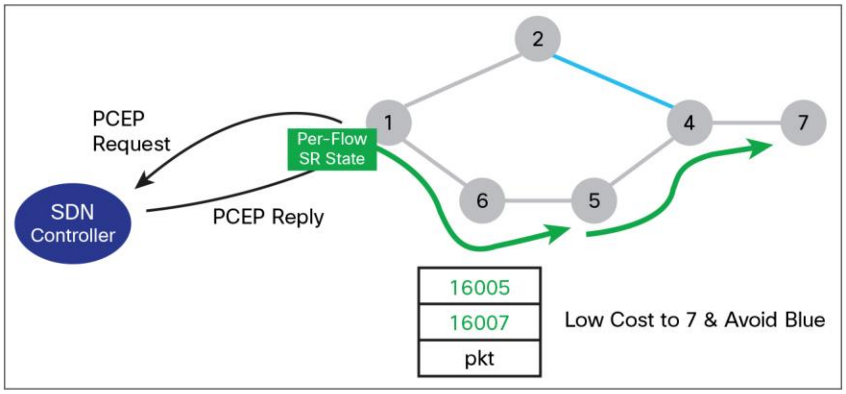

Delivering Differentiated WAN Services on a Per-Application Basis

Frequently, different classes of service and applications need different path characteristics.

In the example shown in Figure 10, a single-area international network with presence in four different regions of the world has lots of cheap network capacity from Tokyo to Brussels through the United States and some scarce, lower-latency, expensive capacity through Russia.

In such a case, the WAN IGP metrics would be tuned to have a shortest path from Tokyo to Brussels through the United States because of available capacity.

This scenario provides efficient capacity planning usage while fulfilling the requirements of most of the applications. However, it may not meet the latency requirements of certain applications, where traversing the lowest-latency path between Tokyo and Brussels is desirable.

In this case, traffic from a server that requires the lowest possible latency should be sent from Tokyo to Brussels through Russia. With segment routing you can easily implement this policy as follows: Router A has a segment-routing traffic-engineering tunnel that sends packets from the applications with the requirement for low latency with the segment stack that first steers the traffic toward Russia. You can achieve this setup by using an anycast SID that is assigned to the routers located in Russia, and then takes the shortest path toward the destination in Brussels.

Figure 10. Delivering Differentiated WAN Services on a Per-Application Basis

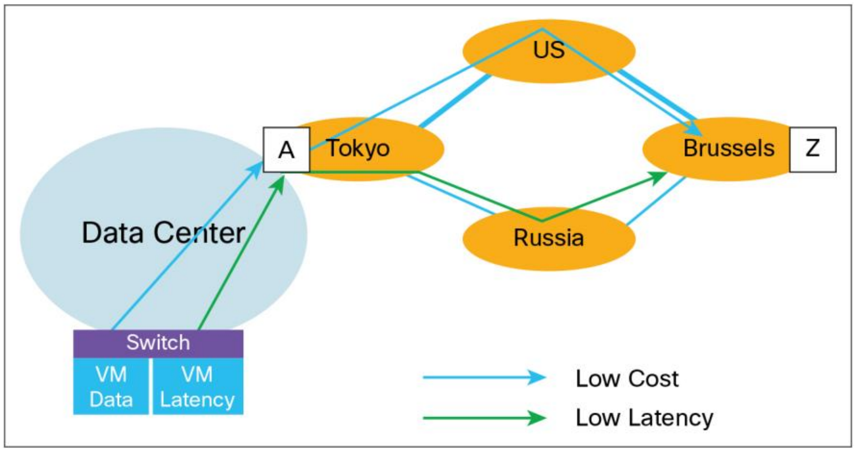

Segment routing brings several benefits to this use case:

- Configuration and state are only on router A, with zero per-service state and signaling at midpoint and tail-end routers.
- ECMP awareness: From Tokyo to Russia traffic travels through all available ECMPs, and the same is true from Russia to Brussels, so all ECMP paths are used.
- Node resiliency property: The traffic-engineering policy is not anchored to a specific core node whose failure could affect the service; instead all (any) routers in Russia are used to fulfill the required policy.

Centralized Egress Peer Engineering

The segment routing-based egress peer engineering use case allows a centralized SDN controller to program any egress peer policy at ingress border routers or at hosts within the segment-routing domain. This capability is extremely valuable for optimizing content delivery.

In this use case a centralized SDN controller uses segment routing to instruct an ingress provider edge or a content source (possibly a server) within the segment-routing domain to use a specific egress router and a specific external interface to reach a particular destination.

Figure 11 illustrates this use case. Autonomous system 1 (AS1) is the only part of the network where segment routing is enabled. Nodes A and B could be top-of-rack switches within a data center, or they could be servers or routers. When sending traffic toward destinations located in AS4, the AS1 operator wants full control of how traffic exits at AS1, and wants very granular and programmable control without dependencies on the routers that exist in the other autonomous systems.

With segment routing, the operator of AS1 can use an SDN controller to program nodes A and B in a way that they will push a segment stack on their packets that will steer the traffic toward router C and specifically toward an interface or peer of router C; for example, A could be programmed to send most of its traffic with segment 100, 101 (blue flow), in which case the traffic will arrive on C through segment 100 and then be forwarded on the link toward D; the C link toward D has segment 101. It's worth noting that this situation is locally significant only for router C, and D does not participate in any segment-routing interaction. Following this same approach the orange and green flows can be programmed by the SDN controller and delivered by the segment-routing-enabled AS1 network.

Figure 11. Segment-Routing Centralized Egress Peer Engineering Use Case

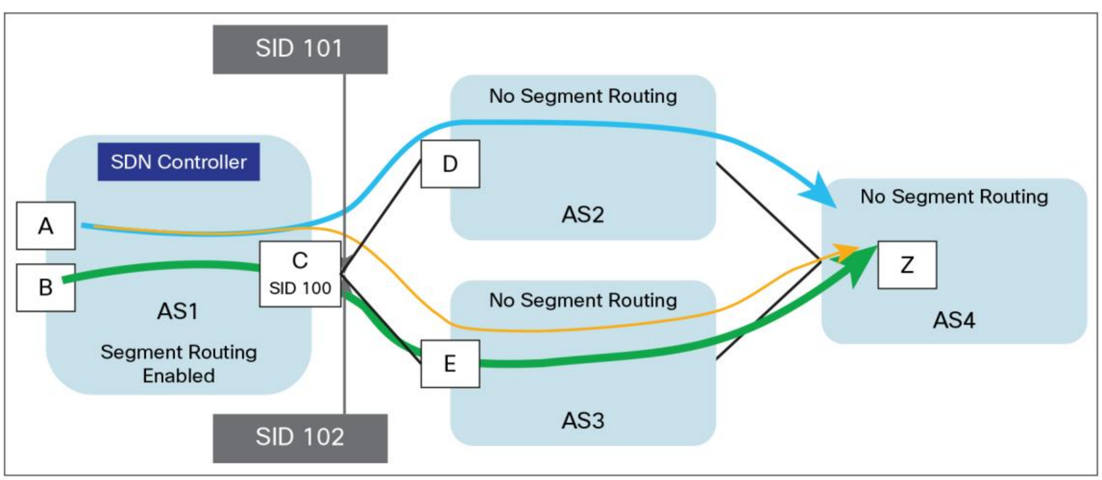

SDN and Segment Routing

Segment routing was built for the SDN era. The combination of a SDN controller with a segment-routing-enabled infrastructure is extremely powerful.

A SDN controller with a global view of the network that can capture business policies and requirements and translate them into segment-routing segments or instructions allows operators an almost infinite number of possibilities for how to provide differentiated services and optimize their network.

Cisco Application Engineered Routing provides a solution that combines an open SDN controller and a segment-routing-enabled infrastructure as defined in Figure 12.

Figure 12. Cisco Application-Engineered Routing

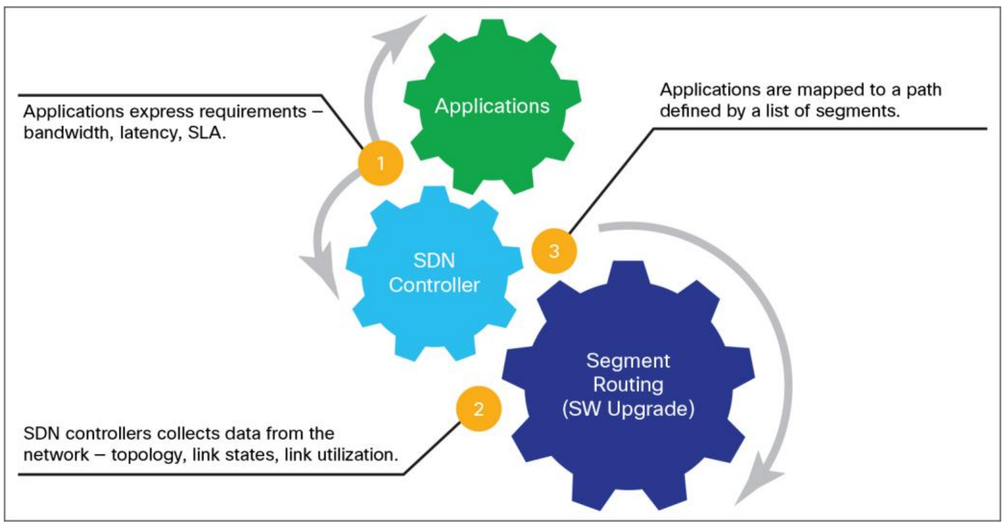

Conclusion

Segment routing enables a unified end-to-end policy-aware network architecture from servers in the data center, through the WAN and up to the MAN and aggregation.

Segment routing is designed for SDN because it seeks the right balance between distributed intelligence, centralized optimization, and application-based policy creation. Other benefits of segment routing are related to operational simplicity, better scale (the segment-routing policy is in the packet), and better use of the installed infrastructure (lower capital expenditures [CapEx]).

Segment routing is being standardized in the IETF, and it enjoys strong industry support.

More Information

Cisco field and partners can test and demonstrate segment routing by accessing the Cisco dCloud (http://dcloud.cisco.com) and searching for segment routing in the demo catalogue.

Cisco customers can request access to segment-routing test setups through their Cisco or partner account teams.

For a consolidated list of IETF drafts and presentations related to segment routing, please visit: http://www.segment-routing.net/.

For any questions, please send an email message to: ask-segment-routing@cisco.com.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-734250-00 07/15