

Segment Routing Overview and Migration Guidelines

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Current Network Deployments](#)

[Overview of Segment Routing](#)

[Why Segment Routing?](#)

[Segment Routing Convergence Benefits](#)

[Segment Routing Control Plane](#)

[Segment Routing Data Plane](#)

[SDN Controller \(SR-PCE\)](#)

[SR Global Block Planning](#)

[Advantages of Homogenous SRGB Block](#)

[Allocate SRGB Block](#)

[Segment Routing Interworking Scenario](#)

[SR to LDP Interworking](#)

[LDP to SR Interworking](#)

[LDP to SR Interworking](#)

[LDP over SR](#)

[Segment Routing Mapping Server](#)

[Segment Routing Migration Guidelines](#)

[Outside-In Strategy](#)

[Inside-out Strategy](#)

[Ship in the Night Strategy](#)

[MPLS LDP Migration to Segment Routing](#)

[RSVP-TE Migration to Segment Routing Policy](#)

[Segment Routing Policy](#)

[Troubleshoot](#)

[Related Information](#)

Introduction

This document describes segment routing migration strategies, which is all about simplifying the transport network and at the same time make it Software Defined Network (SDN) ready. Segment routing is supported with Multi-protocol Label Switching (MPLS) and IPv6 data plane, the main focus for this document to cover migration strategies for MPLS enabled network. This document also highlights the benefits of moving to segment routing and covers some general guidelines to be followed when you plan for migration.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command

Current Network Deployments

MPLS has become the leader and provides various types of Virtual Private Network (VPN) services in the past couple of years. In a very short span of time, MPLS has evolved as a mainstream technology used by Service Provider to create various revenue-generating services like layer 3 VPN, Layer 2 VPN, SLA based services like high bandwidth or low latency path on top of traffic engineering.

Service Provider deployed MPLS with control plane protocols like Label Distribution Protocol (LDP)/BGP for label distribution to achieve traffic forwarding in a service provider domain. Different services offerings like Layer 3 VPN, Layer 2 VPN (point to point vs multipoint) have used MPLS as a transport in a seamless manner. With the demand for meeting specific SLAs for the premium customers, the requirement for traffic engineering become evident and hence Resource reservation Protocol (RSVP) was enhanced to meet this demand. MPLS RSVP Traffic Engineering (TE) had opened up several business use-cases for Service providers like better utilization of available bandwidth, providing low latency path or higher bandwidth to customers.

IP/MPLS networks have become operationally expensive to manage over a period of time due to complex protocol interactions like LDP & IGP Sync, requirements like traffic engineering which was full-filled by RSVP-TE. Network infrastructure and its operations are growing at an exponential pace and becoming more complex. Network owners are looking for a transport technology that can simplify the network by offloading complexity and at the same time open to be programmed via a centralized controller. They are looking for innovative ways to tie the business logic to the underlying network in an efficient & scalable manner, for example, meeting service level agreement (SLA) requirements per application. A technology that can bridge the gap between current network paradigm and futuristic SDN enabled and programmable network.

With the continuous demand and evolution, the MPLS control plane equation has become operationally expensive. As experience is gained from the deployment of this solution, some drawbacks become evident and hence more requirements are added to the goal section, the expectation was to have an improved solution. This iterative process resulted in segment routing evolution.

Overview of Segment Routing

Segment routing is a source-based routing architecture. A node chooses a path and steers a packet through the network via that path by inserting an ordered list of the segment, instructing

how subsequent nodes in the path, that receive the packet, should process it.

Segment routing simplifies operations and reduces resource requirements in the network by removing network state information from intermediary nodes and path information is encoded as an ordered list of segments in label stack at the ingress node. In addition to this, because the shortest-path segment includes all Equal-Cost Multi-Path (ECMP) paths to the related node, SR supports the ECMP nature of IP by design. These two features offer drastic gains in network simplicity & scalability. These gains are achieved by eliminating resource-intensive control plane signaling protocols of MPLS and moves intelligence to the headend device in distributed deployment vs to a centralized controller in a centralized deployment hence reducing the complexity from the network to a greater extent.

Segment Routing can be directly applied on top of MPLS transport with no change on the forwarding plane. The segment to process is on the top of the stack the same as MPLS. Upon completion of a segment, the related label is popped from the stack. Segment Routing is such next-gen technology that can be seamlessly deployed in today's MPLS brownfield network deployment and offer a simple and SDN ready network. The main focus of this document is to describe a migration approach to segment routing for the MPLS data plane.

The SR architecture by design can leverage both distributed and centralized network control model to provide efficient network solutions for Service Provider. The distributed intelligence of the network is used to build these segments at the ingress node, adaptable to any network topology change and pre-calculated backup path against node or link failures that can be activated within sub-milliseconds. The centralized intelligence can focus on network resource optimization by pushing optimum end-to-end paths in the network by a centralized entity. Segment Routing thus allows operators to take advantage of very flexible networking needs for their applications while preserving network resources at the same time.

The integration of segment routing with a centralized controller opens up diverse use-cases and makes the network ready for SDN. Segment routing is good to be deployed in WAN, Access network & Data Centres and an ideal technology for an end to end transport not just limited to Service Provider.

Why Segment Routing?

While data plane in MPLS has rarely been challenged, various control plane protocols for label signaling have added operational complexity as well as poses scalability challenges. To take an example, LDP and its interaction with IGP (LDP-IGP synchronization RFC 5443, RFC6138) have complicated relationships and became an operational challenge for service provider (SP) deployment. On the RSVP-TE side, from the bandwidth reservation point of view, providers who have deployed; have reported it operationally very expensive. As RSVP-TE maintains signaling states on all devices along the path, it has inherent scalability problems. For most of the providers, RSVP-TE was limited to Fast-Reroute (FRR) use-cases.

The table here provides a high-level comparison of RSVP-TE vs SR traffic engineering policy:

RSVP-TE	SR Policy
In the case of RSVP-TE, each path, when computed, needs to be signaled, and state for each path must be maintained in each node traversed by the path. RSVP-TE is used to build a traffic-engineering tunnel, only one path is	Segment routing allows you to implement traffic engineering without a signaling component. Therefore, its architecture scales significantly more, this also simplifies the hardware requirements for the routers in the core of the network (P routers). If ECMPs are present in the network, segment-routing traffic-engineering tunnels can use all the paths for load balancing fl

selected.

across them.

Segment Routing is a promising technology focused on addressing the pain points of existing IP and MPLS networks in terms of simplicity, scalability, and ease of operation. Due to its 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 the operator can build those specific paths based on application requirements.

As mentioned earlier, one of the key characteristics of segment routing is simplicity. These key points summarize this from a different perspective:

- From a configuration standpoint, the number of lines required to enable segment routing is minimum, usually three lines of configuration to make it work.
- From an operational view, it simplifies the operation of an MPLS network by making the label value constant across the core of the network. Hence troubleshooting becomes easier.
- From a futuristic and deployment flexibility perspective, segment routing is especially powerful in the era of SDN. Application requirements program the network; traffic engineering and segregation are done at a much finer granularity (for example, application-specific).

Service providers are looking for more commercial use-cases and exploring ways to make their network infrastructure open to be programmable or SDN ready. SR with a centralized controller makes complete sense here where the controller can further take away the path computation burden off the edge nodes, enabling end-to-end control across multiple domains. Segment routing opens up the potential of a new revenue stream for Service Provider by making the network simpler and SDN capable. 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 Convergence Benefits

With the development of segment routing, link-state IGPs like OSPF & ISIS have been enhanced to distribute segment routing information as well, along with topology and reachability information that they currently signal. In a segment-routing network using the MPLS data plane, segment routing information also known as the segment ID (SID) list is a stack of MPLS labels. 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.

Hence Implementing SR is a low-risk initiative considering that major control plane label distribution protocols and their associated footprints will be offloaded, which will eventually make the network operationally simpler and stable by eliminating the need for protocol interaction.

Another benefit that segment routing brings is automated and native Fast Reroute (FRR) capability or TI-LFA capability, with sub-50-millisecond convergence time. FRR has been deployed in order to cope with link or node failures in a production network. 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 post-convergence path, avoiding transient congestion and suboptimal routing while simplifying the operation and deployment.

Some of the Topology Independent – Loop-Free Alternate (TI-LFA) benefits are:

- Sub-50 msec link, node, and SRLG protection
- 100% coverage on multiple topology scenarios

- Simple to operate and understand
- Automatically computed by the IGP, no additional protocol required
- No state created outside the protecting state at the PLR
- Optimum, the backup path follows the post-convergence path
- Incremental deployment
- Also applies to IP and LDP traffic

Segment routing can be seamlessly deployed in today's MPLS networks as it allows incremental and selective regional 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 data planes.

Segment Routing Control Plane

The control plane of SR defines how the segment ID information is communicated among devices in the network. In the SR network, segment identifiers are advertised via the link-state IGP protocol. Link state IGPs like OSPF & ISIS have been extended to support the distribution of segment IDs. The extensions of IGP protocols would allow any router to maintain a database of all nodes and adjacency segments. Since IGPs are carrying the segment ids, labels in case of MPLS data plane; a separate label distribution protocol is not required as stated previously.

Another element of the control plane of SR deals with how an ingress node is instructed to select the SR path that a packet should follow. There are a couple of ways like static route, distributed vs centralized methods that can be chosen.

Segment Routing Data Plane

The data plane of SR defines how to encode the sequence of segments to be applied on a packet, and how each device should process a packet based on a segment. Defined SR architecture is agnostic to the actual protocol used to carry the information of the SR header in the data plane.

Any router enabled with SR supports the below data-plane operations:

- **CONTINUE** - Forwarding action performed based on active segment.
- **PUSH** - Add a segment ahead of the SR header of the packet and set that segment as the active segment.
- **NEXT** - Mark the next segment as the active segment and execute the instruction encoded by the new active segment.

As stated, 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 Operation LDP Operation

SR Header	Label Stack
Active Segment	Topmost Label
Push Operation	Label Push
Next Operation	Label Pop
Continue Operation	Label Swap

Note: Segment Routing basic building blocks and features can be accessed [here](#).

SDN Controller (SR-PCE)

Software-Defined Networks (SDN) and SDN controller are loaded terms and definition varies. In some cases, these networks are all-encompassing and involve all the topics of orchestration, automation, service assurance, and management of flows within the network. In the following discussion, we touch upon only the flow management component of SDN

The Segment Routing control plane can run purely as a distributed control plane, or it can use a hybrid approach where more complex forwarding paradigms (such as inter-domain routing) are required. The hybrid approach splits the responsibilities: the routers distributed through the network host some functions while external SDN controllers calculate others, for example, the definition of Segment Routing policies and inter-domain paths. In both approaches, the distributed routers run those functions needed to rapidly distribute the link-state database, as well as calculate the shortest path routing tables, monitor the links to the attached nodes, and rapidly recover in the event of a failure.

Segment Routing does not require an external controller function, but as the segment-routing-policy use cases become more complex, or the network increases in scale and extends beyond a single domain, then the use of an SDN controller becomes more important.

Cisco's SDN controller, called Cisco Segment Routing – Path Computation Element (SR-PCE), is based on the Cisco IOS® XR network operating system and can be hosted on either a physical or virtual device. SR-PCE has a northbound interface to the application layer via APIs. Southbound into the transport network, it collects the topology using standards-based protocols such as BGP-LS and subsequently is able to compute and deploy Segment Routing policies across the network. The Segment Routing policy algorithms used by SR-PCE have been purpose-built and specifically designed around segment routing.

For some providers, transport networks will be extremely large and built using multiple domains. In these environments, it is important to isolate the domains as much as possible. At the same time, the operator needs to be able to provide end-to-end services that span domains.

The previous figure shows the solution using a combination of On-Demand Next-hop (ODN), Cisco SR-PCE, and automated steering. This allows an operator to build large complex environments using minimal information exchange between domains, and thus reducing overhead on the network equipment.

When a service needs to span multiple domains, BGP exchanges service routes that have the appropriate SLA identifiers attached. Automated steering then selects the appropriate SR policies while a combination of ODN and SR-PCE builds the multi-domain on-demand Segment Routing policy to the egress device in order to satisfy the service's SLA requirements. Segment routing for traffic engineering (SR-TE) uses a "policy" to steer traffic through the network. Each segment is an end-to-end path from the source to the destination, and instructs the routers in the network to follow the specified path instead of following the shortest path calculated by the IGP or calculated by SR-PCE. If a packet is steered into an SR-TE policy, the SID list is pushed on the packet by the head-end. The rest of the network executes the instructions embedded in the SID list.

SR Global Block Planning

Segment Routing Global Block or SRGB is the range of labels reserved for segment routing when using MPLS as a data plane. This needs to be done on each segment routing aware router in the

network. SRGB is locally significant on a node performing segment routing.

The size of SRGB dictates the number of global segments that can be used in SR deployment. If we go by a typical SP deployment, this relates to the number of routers in the IGP network assuming at least one node segment per router. There could be other prefix segments required for other loopback addresses like Anycast Prefix-SID or prefixes received by redistribution from other parts of the network. Network slicing is another interesting use case where multiple SIDs per node are recommended based on a number of algorithms used.

In Cisco Implementation, The SRGB default block is 16000 to 23999 and It is sufficient for most of the segment routing deployment. It is advisable at the same time to extend this range during the initial planning/deployment phase of SR by keeping the current and future network growth & design use-cases in mind. While It is possible to extend/grow the SRGB size at a later stage, however upfront planning when introducing Segment routing can ensure stable and consistent SRGB which in turn can simplify the network operations. This is also important to avoid traffic flows disruption in the network due to the reconfiguration of this range in the future. It is recommended to use the same SRGB block whether default or non-default SRGB ranges across multiple network domains or nodes within the domain.

Note: In brownfield networks, it is advisable to verify the current label allocation value when you define non-default SRGB range to avoid service disruption.

Advantages of Homogenous SRGB Block

It is strongly recommended to use identical SRGBs on all nodes for homogenous SRGB within the SR Domain. Doing so provides multiple operational and management advantages.

- Use of homogenous SRGB, the MPLS forwarding entries on any router in the network are greatly simplified and it is much easier to correlate them to their IPv4/IPv6 prefix destinations
- By using homogenous SRGB, operations and troubleshooting simplified to a great extent as the same label represents the same global segment at each node.
- Computation of local label value of a prefix SID is straightforward if the SRGB consists of a single consistent range of labels. In such cases, the local label is simply computed by adding the SID index to the SRGB base value.
- The implementation and operation of Anycast-SID become simple and straight-forward when using homogenous SRGB across the network.

Allocate SRGB Block

There are some general guidelines that focus on better manageability to differentiate SID allocation in the network domain.

- Cisco recommends encoding some context like region, country or loopback, etc. into the SID value for loopback0, which will be the node SID for the router in the SR domain.
- It is recommended to select SRGB base values that can easily be mapped and correlate by a human operator (eg. SRGB base is a multiple of 10000) for easier manageability and identification of prefixes.

Segment Routing Interworking Scenario

The MPLS architecture permits concurrent usage of multiple control-plane label distribution protocols such as LDP, RSVP-TE and segment-routing IGP. The control plane of segment routing co-exists with LDP and RSVP is suggested make before the break approach in this article.

The end to end network needs to interwork, means from Segment Routing parts of the network to LDP-only parts of the network and vice versa, end to end MPLS data plane LSP should establish. The interworking functionality takes care of Segment Routing to LDP and LDP to Segment Routing connectivity. It also takes care of interconnecting Segment Routing parts of the network over LDP and interconnecting LDP parts of the network over the Segment Routing domain as described in the subsequent sections.


As the data plane for LDP and Segment routing is label forwarding, this SR/LDP interworking works in a seamless manner. No specific configuration is required to make this work apart from a Mapping server for label assignments to reach LDP only destinations. Traffic forwarding works automatically at any node on the border between LDP and Segment Routing Domain. The seamless interworking is achieved by replacing an incoming label from one protocol by an outgoing label from the other protocol.

These four deployment models are possible and SR-LDP interwork seamlessly:

1. LDP to SR
2. SR to LDP
3. SR over LDP
4. LDP over SR

SR to LDP Interworking

In this deployment model, 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. This is the scenario when LDP is not enabled in the SR domain.




Flow

When a destination is not SR enabled, the SR nodes don't have a prefix-SID for that destination, So no SR transport possible. SR Mapping Server (SRMS) is needed in this case to advertise prefix-SIDs on behalf of non-SR nodes. SR nodes install Mapping Server advertised prefix-SIDs in their forwarding table and establish SR connectivity to non-SR destinations within the SR domain.

LDP to SR Interworking

In this deployment model, 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.




SR Flow

When a node is LDP enabled but its next-hop along the SPT to the destination is not LDP enabled. Any node on the LDP to Segment Routing border (node 3 in this case) will automatically install these LDP-to-SR forwarding entries. Instead of programming an unlabeled entry in the forwarding table, node 3 will automatically connect the LDP Label Switched Path towards node 5, to the Prefix Segment of node 5.

LDP to SR Interworking

Segment routing over LDP (segment routing to LDP followed by LDP to segment routing): At the SR/LDP boundary, the segment-routing prefix segment is mapped to an LDP LSP. At the LDP/SR boundary, the LDP LSP is mapped to a segment-routing prefix segment.




SR over LDP

A mapping server is needed if SR Label Switched Path(s) goes from SR Island and terminates in LDP Island. In the SR island, a prefix-SID is needed to install the Label Switched Path Terminating node is LDP-only. A mapping server advertises a Prefix-SID on behalf of the LDP-only node

LDP over SR

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.



LDP over SR

A mapping server is needed if LDP LSP goes from LDP Island and terminates in SR Island. A prefix-SID is needed in the SR island to install the SR Label Switched Path. The LDP-only nodes can't advertise a prefix-SID. A mapping server advertises a Prefix-SID on behalf of the LDP-only

node

Segment Routing Mapping Server

The objective of the mapping server is to advertise Prefix-to-SID mappings on behalf of other nodes. SID-mappings are advertised on behalf of non-SR-capable nodes. It enables SR-capable nodes to interwork with non-SR-capable LDP nodes.

The Mapping Server functionality in Cisco IOS® XR Segment Routing centrally assigns prefix-SIDs (prefix segment identifiers) for some or all of the known prefixes. The mapping server feature has three primary functions: A router must be able to act as a mapping server, mapping client, or both.

A router functioning as SRMS performs these functions:

- It allows the user to configure SID mapping entries to specify the prefix-SIDs for some or all prefixes. This creates the 'local SID-mapping policy'.
- The local SID-mapping policy contains non-overlapping SID-mapping entries.
- ISIS advertises the local SID-mapping policy in 'SID/Label Binding TLV'.

If IGP receives a prefix-SID from mapping-server and also from another source, IGP uses:

- For local prefixes
 - Use the Prefix-SID configured under an interface.
 - Use the Active SID-mapping policy
- For remote prefixes
 - Use the Prefix-SID attached to the prefix in an IP reachability TLV
 - Use the Active SID-mapping policy

Segment Routing Migration Guidelines

When operators plan to deploy segment routing, they won't have to swap out network hardware. Sometimes it's just the software upgrade to make the network Segment routing capable. For the brownfield environment, Segment Routing can be enabled in current MPLS networks without any rip and replace strategy and as stated previously, can co-exist with LDP/RSVP-TE with no changes to existing control or data plane operation.

The pace of migration to a new technology especially in the brown-field deployment depends on the availability of seamless migration strategies that allows an Operator to migrate from legacy to new technology with minimal or zero impact in the production network. Segment Routing allows an Operator to incrementally upgrade from LDP to SR without breaking any control/data plane for existing traffic.

While migrating the actual production traffic over Segment Routing, it is a common scenario to see a mix of SR-capable and non-SR-capable nodes within the same IGP domain. There are incremental migration strategies available as covered in this guide where some part of the networks is enabled with Segment Routing while the other part is not. With these strategies, some nodes will run as LDP-Only while the others will run as SR-Only nodes. In such cases, as described previously, a Mapping Server is required to advertise the Prefix Segment ID for all non-SR prefixes for an end to end-labeled switched path (LSP).

As stated previously, while considering a migration approach to new technology in a brownfield environment, it is essential to have minimal to zero service disruption. Make before break approach allows to verify the control plane information well before the data plane is updated with new information. In this way, Cisco simplifies your transition from one control plane technology to another. Below are the operational preferences/strategies that can be followed considering the merits of one over the other.

Outside-In Strategy

Service provider network comprises layered architecture consisting of a Core, Aggregation and access networks respectively. In this strategy, Segment Routing migration starts from the access network and then moves towards Pre-aggregation, aggregation and finally into the core segments.

While core consists of big routers routing traffic between various aggregation and access networks. Aggregation is often the service insertion point in the network from where the services start. Access provides the front-haul which connects the cell sites to the network. The traffic is heaviest in the core, heavier in the aggregation and lighter at the access. If this kind of hierarchy is visualized in the form of concentric circles, the innermost circle would form the core, the next one will form the aggregation and the last or the outermost will form the access.

Changes in the access network have operationally minimal exposure, so starting the SR migration from the access network is less risky. Also, the operator gets real experience by the time they move to aggregation/core.

The methodologies for SR migration are based upon the sequence of SR deployment in various segments of the network. When SR deployment is started from the access rings, i.e. from outside and then perforated towards inside aggregation followed by core, the strategy is termed as Outside In strategy. The below figure depicts this methodology of SR deployment.



Outside in strategy

The main highlights of this approach are:

- SR migration is started from the access network.
- Make the aggregation and core SR ready while migrating the Access circles to SR.
- Work gradually into Aggregation and then into Core segments to make the network fully SR-IGP fabric

Why choose Outside In Migration:

- Low Risk: Service disruptions not widespread
- More devices, but divided into manageable islands e.g. Rings


- Allows operators to gain more experience while moving into Aggregation and Core

Inside-out Strategy

In this strategy, SR migration starts from the core network and then it's way towards aggregation and access network.

The lower number of devices provides the benefit to move the core segment to SR quickly and also helps in optimizing the bandwidth which in turn has a higher business impact. Ideally, this approach is recommended for experienced operators as the impact of service disruption will be significant on their customers.

As the name suggests, this approach advocates the SR deployment first in the core of the network. The core network, in most of the operators, comprises the limited number of nodes so the SR migration operation for the core is less and can be completed quickly. However, the approach poses the risk of a huge traffic impact on the core if anything goes bad. The aggregation and access networks are of much greater magnitude and thus they are considered for migration to SR after the core.



Inside Out Strategy

The key steps in the inside out approach are:

- Start SR migration from Core network
- Start making the aggregation and access network ready for SR deployment while the migration work in the core is in progress.
- Work outside into aggregation and then into access segments

Why choose Inside Out Migration:

- High Impact: Operator can make use of BW optimization in the core
- The lower number of devices provides an opportunity to migrate the whole segment relatively quickly.
- Typically for more experienced operators.
- Service disruptions may impact a significant number of customers and services.

Ship in the Night Strategy

This approach enables you to add segment routing into your environment incrementally and phase out your existing transport protocols when you're ready, thus minimizing the service disruption. This approach is recommended for seamless migration.

The segment routing control plane is enabled over the existing LDP network. LDP and Segment routing works independently. In the Cisco implementation, always LDP will be preferred for data forwarding in such cases. In this way, SR can be enabled in a phased manner as per the approach defined previously per network segment.

“**Ship in the night**” approach will also have these advantages.

- Allows for “make before break” verification
- SR control verification before making the switchover
- Segment Routing Control Plane is enabled on existing LDP network
- LDP and SR are kept independent
- SR and LDP PE can interwork seamlessly

Here is the high-level migration plan in order to enable Segment routing and removal of LDP and RSVP protocol. The implementation will be split into three phases.

Phase 1: SR and LDP co-existence by configuring SR and let LDP be the preferred label imposition method.

Phase 2: Preferring SR over LDP as a label imposition method.

Phase 3: Remove LDP and followed by RSVP-TE if configured.

MPLS LDP Migration to Segment Routing

SR Enablement Phase-1

Initial state: All nodes run LDP. RSVP strategy is covered in a later section.

Step 1. Enable Segment routing under the IGP and SID Configuration for each Loopback.

- In no particular order
- Leave default LDP label imposition preference
- Enable TI-LFA for protection configuration for prefixes.

! SRGB Configuration

```
segment-routing
```

```
global-block <SRGB Range>
```

The SRGB default value is 16000 to 23999. The range can be modified based on network size and requirement. Check the **SRGB planning** section for guidelines to define SRGB block.

! ISIS Configuration

```
router isis <Process ID>
```

```
is-type <ISIS Level>
```

```
net <Net ID>
```

```
address-family ipv4 unicast
```

```

microloop avoidance segment-routing

microloop avoidance rib-update-delay <Delay Timer>

`mpls traffic-eng <ISIS Level>

mpls traffic-eng router-id <Router ID>

mpls traffic-eng multicast-intact

segment-routing mpls

interface Loopback0

passive

address-family ipv4 unicast

    prefix-sid <SID ID>

interface <Interface ID>

circuit-type <ISIS level>

point-to-point

address-family ipv4 unicast

    fast-reroute per-prefix

    fast-reroute per-prefix <ISIS level>

    fast-reroute per-prefix tiebreaker < node-protecting | srlg-disjoint > index <priority>

    fast-reroute per-prefix ti-lfa

```

SR prefer command is not configured in this phase.

In the case of multi-domain IGP architecture with BGP LU (RFC 3107), BGP SID should also be configured with the same index value to avoid label conflict.

! BGP SID Configuration

```

Router bgp <AS number>

address-family ipv4 unicast

    network <Loopback0 IP> route-policy <Policy Name>


```

```
route-policy <Policy Name>
```

```
set label-index <SID Index>
```

Step 2. Verify the control Plane on the devices to ensure that LDP imposition continues to be the primary traffic forwarding mechanism. Segment Routing labels are allocated in the control plane by IGP.

This figure represents the state after completion of enablement phase 1 and SR label is generated for all MPLS nodes.



Routing State in Phase 1

SR Enablement Phase-2

Step 1. All Segment routing capable nodes are configured to **prefer SR label** imposition.

- In no particular order but prefer to start from edge nodes.
- Do not remove LDP label imposition.

! ISIS SR prefer Configuration

```
router isis <Process ID>
```


```
address-family ipv4 unicast
```

```
segment-routing mpls prefer
```

There is no change in forwarding plane with SR prefer and LSP would program with SR label

Step 2. Verify the forwarding-plane.

After completion of enablement phase 2, all nodes will have SR prefer for LSP formation and LDP will not be used for LSP formation. This image represents the state when all nodes run SR prefer.



Segment

Routing State in Phase 2

L2 and L3VPN services will continue without any change at this stage.

LDP Removal Phase-3

Step 1. Verify the forwarding-plane with SR.

Step 2. For LDP/RSVP removal from the network, RSVP-TE should be migrated to SR policy (covered in the following section) and LDP based L2 VPN services (VPWS & VPLS) should be BGP based service model.

Step 3. Configure SRMS in order to advertise prefix SIDs on behalf of non-SR nodes within the IGP domain.


! SR mapping server Configuration

```
segment-routing mapping-server
```

```
prefix-sid-map ipv4
```

```
"ip-address/ prefix-length" "first-SID-value" range range
```

Step 4. As a final step, LDP protocols can be removed and the underlay transport network would be SR only. This image depicts the network state after the removal of LDP.



Segment

Routing State in Phase 3

RSVP-TE Migration to Segment Routing Policy



As stated previously, ship in the night approach enables us to add segment routing into the production network incrementally and phase out the transport protocols that already exist when the network operators are ready and thereby minimize the service disruption. This is applicable to RSVP-TE as well.

An RSVP signaled LSP can have a secondary path configured as SR enabled and once the path is up, traffic can switchover to SR signaled LSP via the same tunnel. After this, the RSVP path can be removed from the configuration.

Step 1. Initially, RSVP tunnels are configured on the device.

! RSVP-TE Tunnel LSP

```
interface tunnel-tell
  ipv4 unnumbered Loopback0
  autoroute announce
  !
  destination 6.6.6.6
  path-option 1 explicit name P2-P4-PE6
```



Segment

Routing State in Phase 1

Step 2. On the existing RSVP TE-tunnel, configure a secondary path option with the use of segment routing.


! Secondary path using Segment-Routing

```
interface tunnel-tell
  path-option 2 explicit name P2-P5-PE4 segment-routing
  commit
```

Step 3. Switchover the tunnel to segment routing path option using the command **mpls traffic-engg switchover**.


! Switchover to SR-enabled path

```
mpls traffic-eng switchover tunnel-te 1 path-option 2
```



Segment

Routing State in Phase 2 Step 4. After the successful migration to the SRTE tunnel, it is safe to remove the RSVP path option as shown in the image.



Segment

Routing State in Phase 3

Segment Routing Policy

In Segment Routing, there is a new concept introduced for tunnels, it's called SR-Policy. For moving to segment routing for current tunnels, the SR path can be configured on a legacy TE tunnel interface. However, for any new traffic engineering configuration, its recommended to configure with SR-Policy.

An SR policy path is expressed as a list of segments that specifies the path, called a segment ID (SID) list. Each segment represents an end-to-end path from the source to the destination and instructs the nodes in the network to follow the specified path instead of following the path calculated by the IGP. Once the packet is steered into an SR policy by an automated or manual way, the SID list is pushed on the packet by the ingress node. The rest of the network nodes executes the instructions embedded in the SID list.

Basically, an SR policy is identified as an ordered list (head-end, color, endpoint):

- Head-end – Where the SR policy is instantiated.
- Color – A numerical value that distinguishes between two or more policies to the same node pairs (Head-end – Endpoint). Every policy between the same node pairs requires a unique color value.
- End-point – The destination of the SR policy

In order to configure a local SR policy, you must complete these configurations:

- Create the segment lists
- Create a policy

Segment Routing Policy Configuration:

```
segment-routing
traffic-eng
segment-list name Plist-1
  index 1 mpls label 100101
  index 2 mpls label 100105
!
segment-list name Plist-2
  index 1 mpls label 100201
  index 2 mpls label 100206
!
policy P1
  binding-sid mpls 15001
  color 1 end-point ipv4 6.6.6.6
  candidate-paths
    preference 10
    explicit segment-list Plist-1
      weight 2
    !
    explicit segment-list Plist-2
      weight 2
    !
  !
!
```

A headend can learn different candidate paths of an SR Policy via different available means like via local configuration, via Path Computation Element Communication Protocol (PCEP) or BGP SR-TE. In a distributed control-plane environment, the candidate path is likely to be learned by the headend via local configuration or automated solution such as Cisco NSO. In a centralized control-plane environment, the candidate path is likely to be learned by the headend from the controller via BGP SR-TE or PCEP.

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Related Information

- segment-routing.net
- [Core Fabric Design & Migration](#)
- [Segment Routing Configuration Guide](#)
- [Technical Support & Documentation - Cisco Systems](#)