

Network Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 29 August 2026

M. McBride  
Futurewei  
Y. Liu  
China Mobile  
Z. Li  
Huawei Technologies  
M. Durmus  
E. Erdogan  
Turkcell  
G. Mishra  
Verizon Inc.  
J. Horn  
Cisco  
25 February 2026

SRv6 Deployment Options  
draft-ietf-srv6ops-srv6-deployment-02

Abstract

When deciding to migrate a network from existing data-plane technologies (e.g. MPLS, SR-MPLS or overlay encapsulations such as VXLAN) to SRv6, common questions involve how to perform the migration, how to minimize impact to the existing network and what techniques are available to support a smooth transition. This document presents various deployment and migration options for networks evolving toward SRv6 from prior transport or overlay technologies.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 29 August 2026.

## Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

1. Introduction . . . . .	3
1.1. Requirements Language . . . . .	3
2. Glossary . . . . .	3
3. Gradual vs Direct Evolution . . . . .	4
4. Deployment Options . . . . .	6
4.1. Ships-In-The-Night . . . . .	6
4.1.1. SITN Migration steps . . . . .	7
4.2. Dual Plane . . . . .	8
4.3. Overlay . . . . .	9
4.3.1. VXLAN to SRv6 Migration . . . . .	11
4.4. Interworking . . . . .	12
4.4.1. MPLS over SRv6 . . . . .	13
5. Considerations . . . . .	14
5.1. IPv6 Address Planning . . . . .	14
5.2. BGP in SRv6 Networks . . . . .	15
5.2.1. VPN Service Design . . . . .	16
5.3. Path MTU . . . . .	16
5.4. Inter-AS VPN . . . . .	16
6. IANA Considerations . . . . .	17
7. Security Considerations . . . . .	17
8. Acknowledgement . . . . .	18
9. References . . . . .	18
9.1. Normative References . . . . .	18
9.2. Informative References . . . . .	19
Authors' Addresses . . . . .	19

## 1. Introduction

Segment Routing IPv6 (SRv6) [RFC8986] is a network architecture that leverages IPv6 data plane encapsulation to enable flexible and efficient traffic engineering. It allows for the creation of explicit paths through the network by encoding routing instructions directly into packet headers. Many operators are looking for direction in how to migrate their existing networks to a SRv6 network. It is common for them to have had an IP/MPLS network for over ten years and now ready for a network refresh. Many are convinced it's time to evolve their network to segment routing. And now that SRv6 is mature, they are often planning on that deployment even if currently running SR-MPLS. How to evolve an existing IP/MPLS network to meet the new demands upon a network? Should they run ships in the night (protocol messages coexist being unaware of each other), utilize various tunneling/overlay techniques, use an interworking translation mechanism or other deployment solution? If they are currently running an IP/MPLS network how should they migrate to SRv6? This draft provides various deployment alternatives to help provide guidance to those wanting to migrate their network to SRv6.

SRv6 can be deployed on a typical single-AS network (such as IP backbone network, metro network, mobile transport network, or data center network) or on an E2E network (such as an inter-AS VPN or carrier's carrier network). Before SRv6 is deployed, IPv6 address planning is needed for SID allocation. IGP and BGP designs are then implemented for network nodes, and the corresponding SIDs are advertised for services such as TE and VPN.

[I-D.ietf-srv6ops-problem-summary] provides an overview of the common problems encountered during SRv6 deployment and operation. It provides a foundation for further work, including potential solutions and best practices to navigate deployment. The purpose of this deployment draft is to provide an overview of the various solutions available for SRv6 deployment, particularly in scenarios involving migration from existing transport or overlay technologies to SRv6.

### 1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

## 2. Glossary

MPLS: Multiprotocol Label Switching

RSVP: Resource Reservation Protocol

SR-MPLS: Segment Routing based on MPLS

SRv6: Segment Routing based on IPv6

SRMS: Segment Routing Mapping Server

SITN: Ships in the Night

VTEP: Virtual Tunnel End Point

### 3. Gradual vs Direct Evolution

Migration from an MPLS network to SRv6 constitutes an architectural transition. Two approaches are possible: a gradual evolution via SR-MPLS or a direct migration to SRv6.

A phased approach introduces SR-MPLS (Segment Routing over MPLS) as an intermediate step. SR-MPLS reuses the MPLS data plane while simplifying the control plane through removal of LDP and RSVP-TE. Because it preserves label forwarding, many existing MPLS platforms can support SR-MPLS through software upgrades. At first glance, this appears to reduce migration risk by limiting changes to the data plane. However, introducing SR-MPLS creates an additional architectural stage that must later be migrated again to SRv6. This results in two sequential transitions: first from MPLS to SR-MPLS and then from SR-MPLS to SRv6. Each transition requires planning, testing, operational adaptation and validation. While SR-MPLS may appear incremental, it does not eliminate the need for eventual IPv6 enablement, hardware validation for 128-bit SIDs or operational readiness for SRv6. Instead, it postpones those activities.

In practice, the overall complexity of migrating directly from MPLS to SRv6 is comparable to the cumulative complexity of migrating from MPLS to SR-MPLS and then to SRv6. A direct transition avoids introducing an intermediate architecture that must later be redesigned or retired. Operationally, this reduces duplicated effort in design validation, tooling updates, documentation, training and service migration procedures.

SRv6 does require IPv6 infrastructure support and hardware capable of processing 128-bit SIDs. However, these requirements must ultimately be addressed regardless of whether SR-MPLS is deployed first. Addressing them once, in a single coordinated migration phase, is often operationally simpler than staging multiple control-plane and data-plane transitions.

For networks that are already IPv6-enabled, such as many data center or 5G mobile backhaul deployments, direct migration to SRv6 is typically the most straightforward and future-proof strategy. Even in IPv4-only environments planning IPv6 adoption, it may be operationally more efficient to combine IPv6 deployment and SRv6 introduction into a single transformation program rather than introducing SR-MPLS as an interim solution.

In greenfield deployments where SRv6 is natively supported, direct adoption avoids unnecessary architectural layering. Similarly, if the operations team is already familiar with IPv6 and SRv6 concepts, a direct evolution from MPLS to SRv6 minimizes transitional complexity.

Within this context, several SRv6 transport evolution models can be considered when migrating from traditional MPLS networks or deploying new SRv6-based infrastructures:

1. Ships-in-the-Night (Dual Stack: Independent SRv6 and MPLS): SRv6 and MPLS operate independently within the same network without interaction.
2. Dual-Plane Deployment: A new SRv6 backbone is introduced alongside the existing MPLS infrastructure. New services are activated on SRv6 and MPLS services are migrated over time.
3. SRv6 Overlay (SRv6 over MPLS/IP): SRv6 is deployed as an overlay over the existing transport network, preserving the underlying infrastructure during transition.
4. SRv6 and MPLS Interworking (Coexistence): Interoperability mechanisms enable communication between SRv6 and MPLS domains, allowing phased service migration.

While both gradual and direct strategies are viable, in many environments direct migration from MPLS to SRv6 avoids the operational overhead of an intermediate SR-MPLS phase and results in a cleaner, single-step architectural evolution.

Similar gradual or direct transition considerations apply when migrating from overlay-based data center fabrics (e.g. VXLAN) or other encapsulation technologies toward SRv6. In such environments, operators may choose to introduce SRv6 incrementally alongside existing overlay mechanisms, deploy SRv6 as an underlay replacing the existing transport or directly adopt SRv6-based network virtualization. The migration principles described in this document therefore apply more broadly to networks evolving from either label based or overlay based architectures toward SRv6.

The following diagram depicts the high level options of gradual vs direct evolution to SRv6. An existing MPLS network can first gradually migrate to SR-MPLS before migrating to SRv6 or it can migrate directly to SRv6 and bypass SR-MPLS deployment:

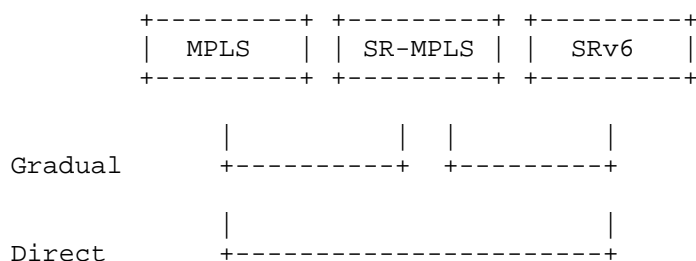


Figure 1: Gradual vs Direct

#### 4. Deployment Options

Various topics are addressed, in this section, to offer options for seamless migration to SRv6. Several SRv6 migration options are highlighted which will each enable a gradual migration and ensures an evolution path without the need for a complete forklift of existing infrastructure.

##### 4.1. Ships-In-The-Night

This solution is a straightforward and popular deployment option. Ships-in-the-Night (SITN) is a technique that allows all routers to run multiple routing processes at once. SRv6 and MPLS operate independently in the same network without interaction. They coexist as separate "ships in the night," with no interworking between them. This technique is commonly used with IPv4 and IPv6 and can also be used with MPLS and SRv6. IPv4 and IPv6 are separate protocols and can't work together without some form of translation mechanism and same is true for MPLS and SRv6. As with MPLS and SRv6, networks run dual stack where both IPv4 and IPv6 run over the same infrastructure as ships-in-the-night.

Ships-in-the-Night is suitable for networks where SRv6 and MPLS serve different purposes (e.g., MPLS for existing VPNs, SRv6 for new services). Complete isolation, of the two control and data planes, avoids interoperability issues and provides flexibility to deploy SRv6 incrementally.

There are drawbacks to running protocols ships-in-the-night such as inefficient resource usage (parallel control planes and data planes) and no synergy between the two technologies. Some routers may struggle with simultaneous MPLS + SRv6. Managing two control planes increases overhead. Some operators prefer gradual migration (Overlay) rather than parallel operation. Maintaining two protocols may introduce additional security vulnerabilities if not managed correctly. Dual-stack networks have an increased attack surface because of both IPv4 and IPv6 being maintained. This may also be true with MPLS and SRv6. The cost of maintaining both networks can be prohibitive as well. Managing and configuring two separate networks can be complex. Ships-in-the-night networks can consume more memory and processing power on networking devices.

The following diagram depicts using Ships-in-the-night SRv6 and MPLS over the same infrastructure:

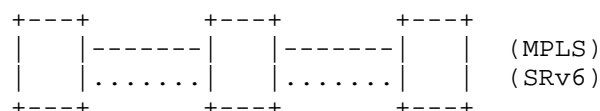


Figure 2: Ships-in-the-night

#### 4.1.1.1. SITN Migration steps

Let's take MPLS L3VPN as an example to describe how L3VPN services can migrate from MPLS to SRv6 using Ships-in-the-night. After network nodes are software upgraded to support SRv6, L3VPN services can be migrated from MPLS to SRv6 using the following procedure:

1. Configure interface IPv6 addresses and locators.
2. Configure IS-IS IPv6 and enable SRv6, and then configure the forwarders to advertise locator routes.
3. Establish BGP peer relationships between the controller and forwarders using the IPv6 unicast address family, and enable BGP-LS and BGP IPv6 SR-Policy. The controller delivers SRv6 Policies, and SRv6 TE tunnels are established on forwarders.
4. On Forwarders, establish BGP VPNv4 peer relationships using IPv6 addresses so that BGP VPNv4 peers advertise VPN routes to each other. The color attribute of the VPN routes is consistent with that of SRv6 Policies to ensure that VPN routes can recurse to the SRv6 Policy.

5. Each forwarder has two routes with the same prefix, one carrying the MPLS VPN label received from the BGP peer established using IPv4 addresses and the other carrying the VPN SID received from the BGP peer established using IPv6 addresses. If the two routes have the same attributes, a forwarder by default preferentially selects the route received from the BGP peer established using IPv4 addresses, and services can still be carried over MPLS tunnels.
6. Configure a route policy so that the forwarder preferentially selects the route received from the BGP peer established using IPv6 addresses. Then, traffic will be automatically switched to SRv6 tunnels, and L3VPN services will be migrated to the SRv6 tunnels.
7. Delete the MPLS tunnel, BGP peer relationships established using the IPv4 unicast address family, and MPLS configurations.

After an SRv6 tunnel is established, and the network is running in SITN mode, services can then be migrated from MPLS to SRv6. Once all services have moved to SRv6, all MPLS related configuration can then be removed.

#### 4.2. Dual Plane

Dual plane refers to building a second, new SRv6-based backbone. The idea is to gradually introduce a separate SRv6 network by first activating new services and later migrating existing MPLS services onto it. A new SRv6 backbone provides a valuable observation window: any protocol bugs, interop issues, or unexpected behavior can be isolated within a limited and controlled environment instead of affecting all existing customers. Once stability is confirmed, new equipment can be integrated into the SRv6 plane, expanding the network organically. If both backbones coexist, an interworking point between the legacy and new domains may become necessary, which could add complexity. With dual plane, the investment cost is also higher, but, when aligned with network renewal cycles, it becomes feasible.

In contrast, the SITN model activates SRv6 directly within the existing network and enables migration in-place. It is conceptually simple and looks appealing, however, in a multi-vendor environment, expecting all routers to run MPLS, SRv6, and other protocols simultaneously without issues can be optimistic. In Türkiye, for instance, frequent fiber cuts also make fast convergence (TI-LFA and RSVP-TE FRR) a real operational challenge. When both mechanisms are active at the same time, every link event triggers recalculations in two separate protection frameworks. This can cause significant load on the control plane.

Once it's confirmed that the various router platforms can handle both planes efficiently, and a dual-plane investment can no longer be justified, then SITN becomes the natural and necessary path forward. Every network and company has their own requirements, and they will move forward by making decisions that best fit these requirements.

#### 4.3. Overlay

With an overlay model, one technology runs on top of the other. The underlying network provides transport, while the overlay provides services. With SRv6 over MPLS, SRv6 packets are encapsulated in MPLS (e.g., in a brownfield migration scenario). SRv6 is deployed as an overlay on top of an existing MPLS transport network. The underlying network remains unchanged, and SRv6 tunnels are encapsulated over the infrastructure. Overlays are useful for gradual migration, allowing operators to introduce SRv6 services without disrupting the existing MPLS/IP core and only minimal changes to the existing network. This allows early adoption of SRv6 features (e.g., network programming, service chaining). There is some overhead due to additional encapsulation (SRv6 headers over MPLS/IP) and it does not fully leverage native SRv6 capabilities in the data plane. It's a common migration technique because migration is fairly easy, it works with existing IPv4 MPLS networks, provides incremental deployment with only the services provider edge (PE) routers needing SRv6 software upgrades. Core network routers can remain IPv4 MPLS (or SR-MPLS) while the rest of the network is migrating to SRv6. How long those core routers remain using MPLS is up to the network operator and can either be a temporary or long term solution depending upon network goals.

For instance, we could utilize a IPv6 provider edge (6PE) overlay if the backbone does not support IPv6. SRv6 services on transit nodes are forwarded through IPv6 over MPLS. 6PE is an MPLS-based overlay mechanism that allows IPv6 traffic to be transported over an IPv4/MPLS core network without requiring IPv6 support on core (P) routers. It leverages MP-BGP and MPLS label stacking to tunnel IPv6 packets across an existing IPv4/MPLS infrastructure. Edge routers connect

IPv6 islands and encapsulate IPv6 in MPLS. When it is challenging to provision dual stack on the core network, a 6PE (or L3VPN, L2VPN, etc) overlay could be used as a temporary migration technique with the capability to evolve to SRv6 in the future. BGP is used to advertise the SRv6 locator and loopback routes of the ingress and egress.

The following diagram depicts using 6PE as the MPLS overlay between SRv6 capable PE nodes:

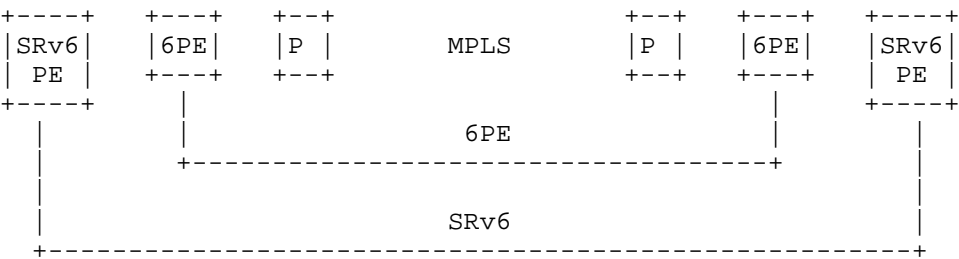


Figure 3: Overlay using 6PE

Overlays can be particularly relevant for multi-vendor networks where some of the multi vendor platforms do not yet support SRv6 or there are other readiness gaps. They may have initiated gradual hardware replacement plans but it is not always possible to invest in SRv6-capable hardware across all vendors and network layers at the same time. For this reason, the overlay approach can be used as a transitional mechanism for operators who want to gain early experience with SRv6 within limited domains during their migration.

Turkcell's network architecture, for instance, uses a layered design, and each layer includes devices from different vendors. In the data center (DC) network, they are using one vendors equipment which will carry the first SRv6 deployment. This will allow them to observe SRv6 behavior directly in a live environment. By starting with a single-vendor domain, they will also have the opportunity to experience the operational simplicity of a homogeneous environment, which will help better understand the added complexity that comes with multi-vendor SRv6 deployments in later phases. In the mobile traffic layers, two different vendors equipment are used together, and these domains include complex L3VPN-based service chaining. These cases are being analyzed separately to assess SRv6 readiness and migration feasibility.

The overlay model is typically not considered a long-term migration path, but rather a transitional deployment approach that provides flexibility during the migration phase. While Overlay models may offer short-term practical advantages, they do not fully leverage native SRv6 data-plane capabilities and may introduce additional encapsulation overhead. For long-term migration goals, Ships-in-the-Night and/or Dual Plane models are typically preferred.

#### 4.3.1. VXLAN to SRv6 Migration

##### 4.3.1.1. Overview

Some data center and metro networks use VXLAN as an overlay to provide L2/L3 services over an IP or MPLS network. As operators look to migrate these networks to SRv6, a clear path is needed to transition VXLAN-based deployments to a SRv6 infrastructure.

##### 4.3.1.2. Overlay Approach

VXLAN packets are encapsulated within SRv6 headers while the underlying IP network remains unchanged. This allows services carried over VXLAN to continue operating while gradually introducing SRv6 capabilities at the edges and on SRv6-capable transit nodes. This is similar in principle to SRv6 over MPLS (6oM) overlays.

Existing VXLAN Virtual Tunnel Endpoints (VTEPs) can be upgraded to support SRv6 encapsulation for north-south traffic, while east-west traffic within a data center may continue using native VXLAN encapsulation during transition.

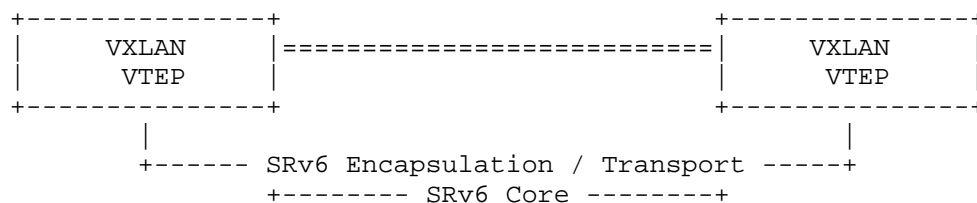


Figure 4: VXLAN over SRv6

##### 4.3.1.3. Migration Steps

The following steps provide a phased approach to migrate VXLAN networks to SRv6:

1. Upgrade edge VTEPs to support IPv6 and SRv6 encapsulation while continuing VXLAN forwarding.

2. Plan and assign SRv6 locators and SIDs for all ingress and egress VTEPs and SRv6 transit nodes.
3. Deploy SRv6 core infrastructure, either natively or as an overlay on the existing IP/MPLS network.
4. Enable SRv6 encapsulation on upgraded VTEPs and steer selected VXLAN traffic over SRv6 transport paths.
5. Gradually migrate VXLAN traffic to SRv6 paths by modifying VTEP policies to encapsulate traffic in SRv6.
6. Validate traffic flows, TE policies and service SLAs.
7. Decommission VXLAN encapsulation once all services have transitioned to native SRv6.

#### 4.4. Interworking

Another migration strategy is to allow an existing MPLS network to interwork with SRv6, rather than only run ships-in-the-night or overlay. [I-D.ietf-spring-srv6-mpls-interworking] describes SRv6 and MPLS/SR-MPLS interworking procedures which can roughly be compared to translation solutions such as NAT or 464XLAT. This strategy enables interworking between SRv6 and MPLS domains in situations where completely separate domains must be maintained. Translation mechanisms (e.g., Segment Routing Mapping Server or SRMS) are used to map SRv6 SIDs between the two domains. This option allows hybrid operation (e.g., SRv6 at the edge, MPLS in the core). Interworking requires additional control-plane mechanisms for SID translation and may add complexity in managing two different forwarding paradigms. New SRv6 behaviors, and MPLS labels, stitch the end to end path across different data planes. The interworking document assumes SR-MPLS-IPv4 for MPLS domains but the design equally works for SR-MPLS-IPv6, LDP-IPv4/IPv6 and RSVP-TE-MPLS label binding protocols. It provides transport interworking solutions such as SRv6 over MPLS (6oM) and MPLS over SRv6 (Mo6) along with service interworking solutions such as SRv6 to MPLS(6toM) and MPLS to SRv6 (Mto6).

Using a gateway is an Interworking (IW) example which supports both BGP SRv6 based L2/L3 services and BGP MPLS based L2/L3 services for a service instance. It terminates service encapsulation and performs L2/L3 destination lookup in a service instance:

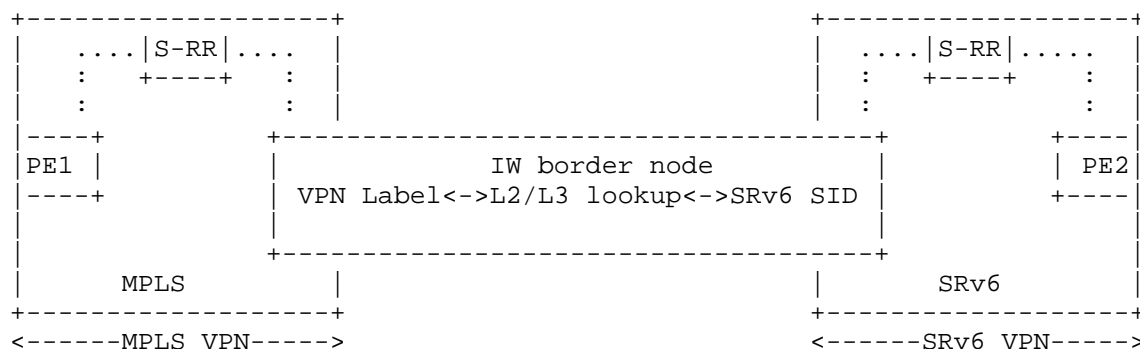


Figure 5: Gateway IW

#### 4.4.1. MPLS over SRv6

In interworking scenarios where the core network has migrated to SRv6, but the access or aggregation layers continue to operate using MPLS, the MPLS-over-SRv6 (Mo6) technology ([I-D.ietf-spring-srv6-mpls-interworking]) can be used to provide seamless service continuity. This approach is particularly relevant for large-scale networks that use BGP-LU to achieve end-to-end MPLS LSPs.

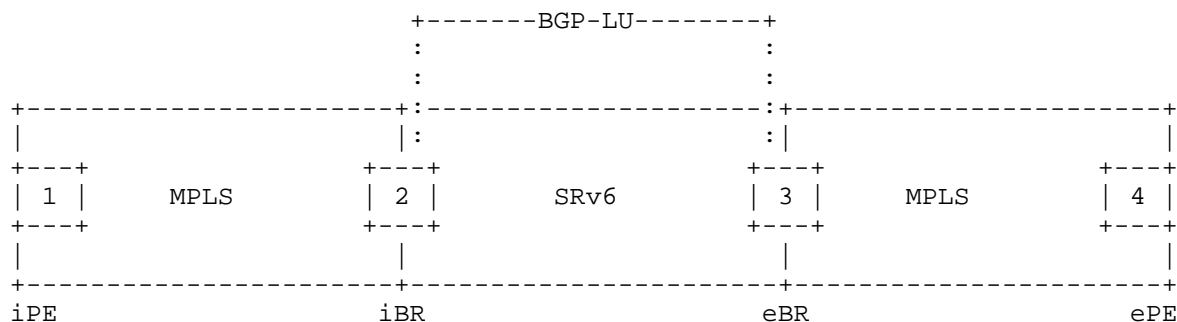


Figure 6: Example of MPLS over SRv6 Interworking

The ingress and egress Border Routers (BRs) perform the interworking between MPLS and SRv6 domains. The BRs exchange loopback prefixes using BGP-LU SAFI, where the SRv6 SID associated with each prefix is an END.DT or END.DTM SID. The prefix may be learned directly via BGP-LU or redistributed from the IGP.

This principle applies equally to traditional MPLS networks that use LDP or RSVP-TE signaling, as well as to networks using SR-MPLS. In each case, the Mo6 mechanism allows MPLS-based transport services to be extended seamlessly across an SRv6 core, facilitating a phased migration strategy while preserving end-to-end service continuity.

## 5. Considerations

Here are a few additional considerations when migration from MPLS to SRv6.

### 5.1. IPv6 Address Planning

SRv6 requires a network running IPv6 and forwards packets based on native IPv6. Interface IPv6 addresses need to be configured prior to SRv6 configuration. IP address planning is an important part of network design and directly affects subsequent routing, tunnel, and security designs. Well-designed IP address planning makes service provisioning and network OAM much easier. When SRv6 needs to be deployed on a network, if IPv6 has been deployed and IPv6 addresses have been planned, the original IPv6 address planning does not need to be modified, and we only need to select a reserved network prefix and use it to allocate SRv6 locators. If neither IPv6 has been deployed on a network, nor IPv6 addresses have planned, IPv6 address planning can be performed by determining the principles for IPv6 address planning on the network, determining the method of IPv6 address allocation, and hierarchically allocating IPv6 addresses.

During IPv6 address planning, for an E2E SRv6 network for instance, each network domain is configured with a network prefix for locator allocations to devices in this domain, allowing advertisement of only an aggregated locator route to devices outside the domain. If no IPv6 loopback interface has been configured on the network, the locator and loopback address with the same network prefix can be allocated so that only the aggregated route shared by the locator and loopback address needs to be advertised, thereby reducing the number of routes. A separate network prefix is allocated to the access and aggregation layers, and another separate network prefix is allocated to the IP core layer. Only an aggregated IPv6 route (locator and loopback address) is advertised between the aggregation and IP core layers. SRv6 service nodes only need to learn the aggregated route and the specific routes in the local domain to carry E2E SRv6 services. In addition, the number of service configuration points is reduced to two: ingress and egress. As such, the specific routes of a domain are not flooded to other domains. In addition, route changes, such as route flapping, in one domain do not cause frequent route changes in another domain. This enhances security and stability within the network.

Operators should consider the guidance in [RFC9602], which updates the IPv6 addressing architecture to describe the use of SIDs as IPv6 addresses. RFC9602 allocates a dedicated IPv6 prefix for SRv6 SIDs and clarifies their structure and semantics. Using the dedicated SRv6 SID prefix can simplify address planning, improve operational consistency, and provide a clearer distinction between infrastructure addresses and SRv6 locator space.

## 5.2. BGP in SRv6 Networks

On an SRv6 network, in addition to the conventional route advertisement function, BGP also supports information exchange between forwarders and a controller. Forwarders use BGP-LS (Link State) to report information, such as the network topology and latency, to the controller for path computation. To support SR, forwarders need to report SR information to the controller through BGP-LS ([I-D.ietf-idr-bgp-ls-sr-policy]). Additionally, the controller uses BGP SR Policy ([I-D.ietf-idr-sr-policy-safil]) to deliver SR path information. For this reason, on an SRv6 network, BGP design needs to consider not only the IPv6 unicast address family peer design and VPN/EVPN address family peer design, but also the BGP-LS address family peer design and BGP IPv6 SR-Policy address family peer design.

In a VPN network (which uses MP-BGP to distribute VPN routes), a Route Reflector (RR) eliminates the need for a full mesh by allowing PE routers to peer only with the RR, which then reflects VPN routes to all other PEs. BGP treats VPNv4 (IPv4 VPN) and VPNv6 (IPv6 VPN) as different address families. Both VPNv4 and VPNv6 need to be enabled in MP-BGP when using both address families in, for example, Ships-in-the-night deployments. A single VPN can be supported by both MPLS and SRv6 simultaneously in SITN mode, but the two control planes operate independently, and seamless interworking requires additional mechanisms. VPN service over SRv6 is described in [RFC9252].

BGP information types have various roles in SRv6. VPNv6 routes carry customer VPN routes with SRv6 SIDs (End.DT6, End.DX4, etc.). BGP-LS collects and distributes SRv6 topology info to controllers (e.g., for SDN) and BGP SRv6 policies distribute SRv6 Traffic Engineering (TE) policies (e.g., Flex-Algo, explicit paths).

#### 5.2.1. VPN Service Design

SRv6 VPN services can use BGP as the unified signaling control plane to provide L2/3 service connections. EVPN can be used to carry both L3VPN and L2VPN services in SRv6, thereby simplifying protocols. Hierarchical VPN is widely deployed on MPLS networks to reduce the number of routes on access devices at network edges. E2E VPN is recommended for SRv6 networks because only service access points, instead of transit nodes, need to be configured. Also, transit nodes do not need to be aware of services, and this in turn facilitates both deployment and maintenance.

#### 5.3. Path MTU

SRv6 encapsulation introduces additional IPv6 header and SRH overhead. In VPN deployments, where multiple encapsulations (e.g., IPv6 + SRH + VPN service headers) may be present, packets are more likely to exceed the default IPv6 Path MTU (PMTU). Exceeding the PMTU can result in fragmentation or packet drops if PMTU discovery is not functioning reliably.

Operators could explicitly account for SRv6 overhead in access and core MTU planning. Common practices include configuring consistent MTU values across the SRv6 domain, enabling IPv6 PMTU Discovery [RFC8201], and reserving sufficient headroom for SRH and VPN encapsulation. During migration or mixed MPLS/SRv6 deployments, operators should validate MTU consistency end-to-end to avoid service interruption.

To mitigate the impact of PMTU variations on live traffic during deployment, operators can use staged rollout and verification procedures. This may include proactive measurement of end-to-end MTU across VPN sites, testing representative traffic flows with encapsulation enabled, and validating that ICMP messages are properly propagated. Where PMTU discovery cannot be assured, setting a conservative maximum packet size at ingress PEs can prevent customer traffic from exceeding the supported path MTU.

#### 5.4. Inter-AS VPN

Inter-AS VPN is widely deployed in MPLS networks and remains critical during SRv6 migration. In SRv6, inter-AS VPN can be realized by extending VPNv6 routes with SRv6 SIDs across ASes using MP-BGP. Depending on the migration strategy, different options can be applied:

With ships-in-the-night, each AS can independently operate MPLS or SRv6 VPNs, with traffic exchanged over dual-stack BGP sessions.

In an overlay model, SRv6 traffic between ASes can be tunneled over existing MPLS or IP interconnects until both domains natively support SRv6.

With interworking, SRv6 SIDs may be translated to MPLS labels (or vice versa) at the ASBR, enabling hybrid deployments while preserving existing inter-AS VPN services.

Operators need to consider the impact on route scaling, locator design, and policy enforcement at AS boundaries. Security measures described in [RFC8754] also apply to inter-AS SRv6 deployments, with additional need to enforce filtering and validation at ASBRs. The procedures for VPN service over SRv6 are further described in [RFC9252].

## 6. IANA Considerations

N/A

## 7. Security Considerations

The security considerations for Segment Routing are discussed in [RFC8402]. Section 5 of [RFC8754] describes the SR Deployment Model and the requirements for securing the SR Domain. The security considerations of [RFC8754] also cover topics such as attack vectors and their mitigation mechanisms that also apply the behaviors introduced in this document. Together, they describe the required security mechanisms that allow establishment of an SR domain of trust. Having such a well-defined trust boundary is necessary in order to operate SRv6-based services for internal traffic while preventing any external traffic from accessing or exploiting the SRv6-based services. Care and rigor in IPv6 address allocation for use for SRv6 SID allocations and network infrastructure addresses, as distinct from IPv6 addresses allocated for end users and systems (as illustrated in Section 5.1 of [RFC8754], can provide the clear distinction between internal and external address space that is required to maintain the integrity and security of the SRv6 Domain. Additionally, [RFC8754] defines a Hashed Message Authentication Code (HMAC) TLV permitting SR Segment Endpoint Nodes in the SR domain to verify that the SRH applied to a packet was selected by an authorized party and to ensure that the segment list is not modified after generation, regardless of the number of segments in the segment list. When enabled by local configuration, HMAC processing occurs at the beginning of SRH processing as defined in Section 2.1.2.1 of [RFC8754].

## 8. Acknowledgement

Thank you to Dhruv Dhody for providing extensive comments on this draft. We also recognize the comments from Dongjie, Yanrong, Liuyao, Nat Kao, Eduard Metz, Cheng Li, Luis Miguel Contreras Murillo and Nick Morris.

## 9. References

### 9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <<https://www.rfc-editor.org/info/rfc8201>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.
- [RFC9602] Krishnan, S., "Segment Routing over IPv6 (SRv6) Segment Identifiers in the IPv6 Addressing Architecture", RFC 9602, DOI 10.17487/RFC9602, October 2024, <<https://www.rfc-editor.org/info/rfc9602>>.

## 9.2. Informative References

### [I-D.ietf-idr-bgp-ls-sr-policy]

Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies using BGP Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-17, 6 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-17>>.

### [I-D.ietf-idr-sr-policy-safi]

Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>>.

### [I-D.ietf-spring-srv6-mpls-interworking]

Agrawal, S., Filsfils, C., Voyer, D., Dawra, G., Li, Z., and S. Hegde, "SRv6 and MPLS interworking", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-mpls-interworking-02, 7 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-mpls-interworking-02>>.

### [I-D.ietf-srv6ops-problem-summary]

Liu, Y., Graf, T., Miklos, Z., Contreras, L. M., and N. Leymann, "SRv6 Deployment and Operation Problem Summary", Work in Progress, Internet-Draft, draft-ietf-srv6ops-problem-summary-00, 22 November 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-srv6ops-problem-summary-00>>.

## Authors' Addresses

Mike McBride  
Futurewei  
Email: [mmcbride7@gmail.com](mailto:mmcbride7@gmail.com)

Yisong Liu  
China Mobile  
Email: [liuyisong@chinamobile.com](mailto:liuyisong@chinamobile.com)

Zhenbin Li  
Huawei Technologies

Email: lizhenbin@huawei.com

Mehmet Durmus  
Turkcell  
Email: mehmet.durmus@turkcell.com.tr

Ersin Erdogan  
Turkcell  
Email: ersin.erdogan@turkcell.com.tr

Gyan S. Mishra  
Verizon Inc.  
Email: gyan.s.mishra@verizon.com

Jakub Horn  
Cisco  
Email: jakuhorn@cisco.com