Network Working Group
Internet-Draft
Intended status: Informational
Expires: 25 March 2026

Y. Liu
China Mobile
T. Graf
Swisscom
Z. Miklos
MTN
L. Contreras
Telefonica
N. Leymann
Deutsche Telekom

25 September 2025

SRv6 Deployment and Operation Problem Summary
draft-liu-srv6ops-problem-summary-06

Abstract

This document aims to provide a concise overview of the common problems encountered during SRv6 deployment and operation, which provides foundations for further work, including for example of potential solutions and best practices to navigate deployment.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on March 25, 2026.

Table of Contents

1. Introduction

Segment Routing over IPv6 (SRv6) is a new technology that builds upon the existing IPv6 infrastructure to offer programmable data plane capabilities. This allows for more granular control over traffic forwarding, enabling flexible and scalable network designs. While SRv6 presents numerous potential benefits, such as improved traffic engineering, optimized resource utilization, its deployment and operation come with certain challenges.

This document aims to provide a concise overview of the common problems encountered during SRv6 deployments and operations, which provides foundations for further work, including for example potential solutions and best practices to navigate deployment . By understanding these challenges and exploring mitigation strategies, network administrators can make informed decisions when implementing and managing SRv6 networks.

This document identifies a number of Deployment and Operation Problems (DOPs) that require additional work within IETF.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. SRv6 Network Migration

2.1. SRv6 Migration from MPLS Networks

In the evolution from non-SRv6 networks to SRv6 networks, the migration from MPLS to SRv6 represents the most typical and common scenario. This process requires addressing the following key challenges:

- Ensuring interoperability between MPLS and SRv6 during the transition.
- Avoiding service interruptions to existing VPN services.
- Supporting smooth migration paths with minimal deployment and operational complexity.

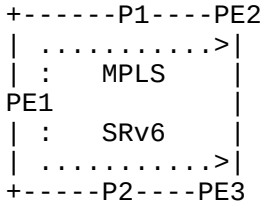

Figure 1: Intra-domain SRv6&MPLS Coexistence

SRv6 and MPLS Coexistence means a network that supports both SRv6 and MPLS in a given domain. This may be a transient state when brownfield MPLS network upgrades to SRv6 or permanent state when some devices are not capable of SRv6 but supports native IPv6 and MPLS. A smooth transition from an MPLS network to an SRv6 network is required. For instance, deploy dual-stack tunnels for VPN over MPLS and VPN over SRv6, with MPLS and SRv6 sharing VPN instances. When the next hop of the route is an IPv4 address, iterate through the MPLS tunnel; when the next hop is an IPv6 address, iterate through the SRv6 tunnel. Prefer VPN routes based on SRv6. Once the transition is complete, the MPLS tunnel can be removed. When operating both MPLS and SRv6 concurrently, key considerations arise regarding how to ensure effective protection during failures, as well as how to manage the increased complexity of performance monitoring and optimization deployment. So it becomes critical to address methods for reducing the complexities associated with deployment and operational maintenance.

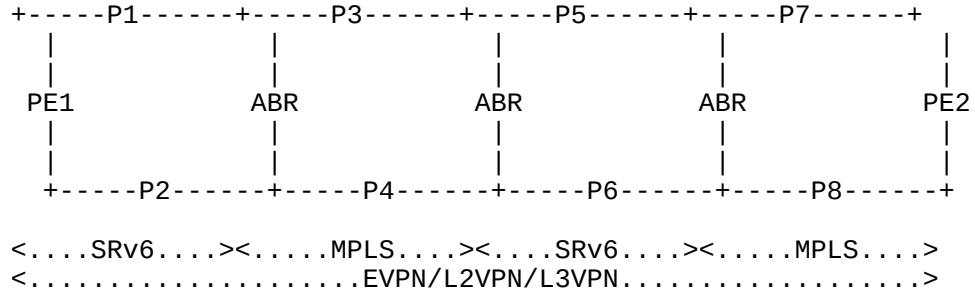

Figure 2: Inter-domain SRv6&MPLS Coexistence

Ensuring seamless interworking between legacy MPLS networks and newly deployed SRv6 networks in long-term coexistence scenarios presents significant challenges, particularly in multi-domain architectures where each domain operates independent IGP instances and employs a single data plane type for both overlay and underlay. The potential cascading of MPLS and SRv6 domains introduces complexities in guaranteeing end-to-end path quality requirements such as latency, bandwidth, and reliability when traffic traverses these heterogeneous data planes. Additionally, achieving rapid end-to-end path convergence during failures requires robust mechanisms, especially when edge nodes must perform route regeneration and re-advertisement functions between different address families like EVPN and VPNv4 or VPNv6, while minimizing operational complexity and maintaining consistent service levels across such hybrid infrastructures remains a critical operational concern.

DOP-1: How to smoothly migrate from existing MPLS networks to SRv6 while ensuring service continuity, interoperability and minimizing deployment complexity.

2.2. SRv6 Inter-Domain Connectivity

In some cases, SRv6 networks need to extend across multiple domains, including third-party or legacy networks that may not natively support SRv6 or even IPv6. This inter-domain scenario introduces new requirements and challenges:

Ensuring SRv6-based services can traverse domains where native SRv6 or IPv6 is not supported.
Reducing complexity compared to traditional MPLS-based inter-domain solutions.
Improving scalability and operational simplicity for service providers.

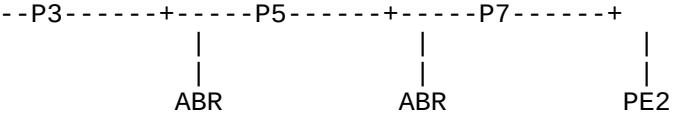

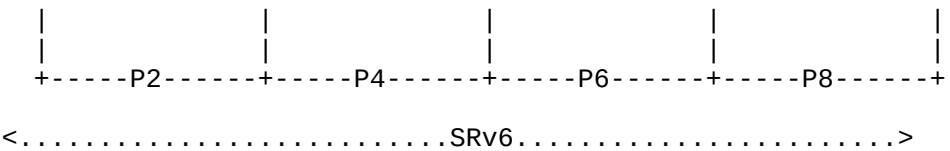

Figure 3: Inter-domain SRv6 End to End

When migrating from BGP/MPLS VPN [RFC4364] inter-domain solutions (Option A/B/C) to SRv6-based architectures, several operational and technical challenges emerge. For Option A, back-to-back VRFs require complex interface-level configurations that conflict with end-to-end service paradigm of SRv6. For Option B, the stateful inter-AS label mapping mechanisms become redundant when transitioning to the source routing model of SRv6, creating protocol coexistence issues during migration. For Option C, recursive label stacking proves incompatible with SRv6 end to end service, necessitating Autonomous System Border Router(ASBR) functionality redesign. Additional challenges include maintaining service continuity during phased migration, retraining operational staff for SRv6 troubleshooting, and achieving consistent traffic engineering across hybrid MPLS-SRv6 domains. The IPv6 infrastructure readiness gap-particularly regarding MTU management and ICMPv6 processing-further complicates deployment, while the OAM mechanisms for inter-domain SRv6 operations in fault detection and performance monitoring are also facing challenges.

DOP-2: How to achieve scalable and simplified end-to-end inter-domain communication using SRv6, overcoming the limitations of traditional MPLS-based solutions.

3. SRv6 Network Visualization

3.1. SRv6 Path Performance Visibility

The existing IETF data collection frameworks can be applied to SRv6 for both data plane and control plane monitoring, but currently lack critical capabilities for measuring SRv6-specific path performance metrics and granular traffic statistics. This significant visibility gap fundamentally limits the ability to conduct meaningful SRv6 network analysis, either making it completely impossible or severely compromising its effectiveness-particularly for use cases such as:

Network delay and packet loss measurement: In scenarios where an SRv6 path traverses multiple network segments (e.g., edge, aggregation, core network), current monitoring tools are unable to provide accurate per network segment and end-to-end delay or packet loss data.

End-to-end path reconstruction: Traditional network monitoring only provides hop-by-hop metrics visualization, which offers limited capabilities for end-to-end SRv6 path optimization and adjustment. Simply combining per-hop metrics does not accurately reflect the actual performance of the overall path, and thus fails to provide effective support for SRv6 path re-optimization.

These issues make service path validation, traffic forecasting, and microsecond-level troubleshooting challenging in SRv6 networks.

DOP-3: The collection of SRv6-specific path performance data is incomplete and inefficient, limiting end-to-end visibility of SRv6 service paths and making it difficult to validate and optimize performance.

3.2. Multi-source Data Troubleshooting

In network fault management, a key challenge lies in the inability to automatically correlate information from multiple monitoring sources such as BGP Monitoring Protocol(BMP), Internet Protocol Flow Information Export(IPFIX), and YANG-push. Currently, when a failure occurs, operators only can manually collect and interpret data from these isolated systems to identify the root cause. For example, in the case of an SRv6 service interruption:

BMP can report report routing instability or BGP session changes;
IPFIX could indicate abnormal traffic drops for specific SRv6 paths or segments;
YANG-push can show abnormal resource utilization or interface errors from device telemetry statistics.

Although each traffic provides valuable insights, the absence of a unified correlation mechanism or common data model requires manual cross-referencing and time-consuming analysis. This lack of

automated correlation significantly delays fault detection, diagnosis, and service recovery. Furthermore, the diversity of data models and semantic differences in SRv6-related telemetry create additional integration barriers.

Most current tools are unable to effectively aggregate, interpret, and present multi-source SRv6 data in a consistent and actionable manner, which further impacts operational efficiency and network reliability.

DOP-4 Multi-source network data (BMP, IPFIX, YANG, etc.) lacks automatic correlation and integration, complicating SRv6 network operation and fault analysis, and leading to delays in troubleshooting and recovery.

4. SRv6 Address Planning

Existing IPv6 address planning are primarily based on network types and hierarchical administrative divisions. While effective for traditional IPv6 deployment, such approaches are insufficient to meet the requirements of SRv6 Segment Identifier(SID) allocation, especially in the context of advanced capabilities such as SRv6 compression. If SRv6 SID planning simply inherits the conventional IPv6 structure, it may lead to a fragmented SID space, complicating end-to-end segment routing. On the other hand, deviating entirely from the existing addressing scheme introduces significant complexities in address management and operational consistency.

In multi-domain networks, high levels of route aggregation are desirable for efficient SRv6 compression. However, this objective often conflicts with existing IPv6 address planning that are organized along administrative boundaries. Consequently, a rethinking of SRv6 address planning is necessary to align compression benefits with scalable network design. Strategic allocation of SRv6 SID blocks must holistically consider both administrative division management and route aggregation requirements. Similarly, the distribution of NodeIDs and functions should balance administrative practicality with optimal address space utilization within the SRv6 architecture.

Some initial approaches and considerations for structured SRv6 SID assignment can be found in [I-D.liu-srv6ops-sid-address-assignment], which provides a foundation for further standardization and operational practices.

DOP-4: Exisiting IPv6 address planning methods are insufficient to accommodate the structural requirements of SRv6 SIDs. The inherent complexity introduced by the SID architecture extends beyond

conventional IPv6 addressing, complicating overall network planning and integration.

DOP-5: In inter-domain deployment environments, SRv6 SID allocation poses challenges such as inefficient utilization of address space, impediments to route aggregation, and inconsistent compression performance, which may undermine the efficiency gains promised by SRv6.

5. Traffic steering to SRv6

The general purpose of traffic steering is to optimize the allocation and transmission of network resources, ensure a balanced distribution of network traffic, improve network performance, reduce congestion, and increase available bandwidth to provide users with a better network experience.

In SRv6-enabled networks, traffic steering plays a critical role, especially in realizing advanced use cases such as service function chaining, traffic engineering, and end-to-end SLA assurance. However, steering traffic to SRv6 paths introduces unique challenges, as SRv6 supports a wide range of encapsulation and policy mechanisms, and these choices greatly affect deployment flexibility, operational complexity, and optimization capabilities.

There are various SRv6 traffic steering methods, each with its own unique advantages and limitations. For example:

- Using BGP color community-based policy may fail to provide sufficient granularity or flexibility for dynamic adjustment in some enterprise scenarios.
- Using flow-based steering such as flowspec may become infeasible when facing a large number of flows, as maintaining per-flow granularity overwhelms the control plane and devices.

Therefore, it is essential to select the appropriate SRv6 traffic steering mechanism based on the specific application scenario. Some initial technical considerations can refer to [I-D.geng-srv6ops-traffic-steering-to-srv6].

When selecting network traffic steering methods, factors such as network architecture, service requirements, resource constraints,

and operational costs must be comprehensively considered, and the selection logic varies significantly across different hierarchy nodes. For example, delay-sensitive and bandwidth-intensive scenarios have different requirements for traffic steering methods. Moreover, strategies for selecting traffic steering methods differ depending on network scale. Finally, operational complexity varies with different steering methods, influencing operational requirements.

DOP-6 There are various methods for SRv6 traffic steering, making it difficult to select the appropriate method for different scenarios. This leads to deployment complexity, especially when the chosen method does not meet the requirements of granularity, scalability, or operational ease. For example, using a simple color-based policy may not support fine-grained tuning, and flowspec-based approaches may not scale in high-flow-volume environments.

6. Deployment Practice for SRv6 Protection

Implementing reliability practices can significantly enhance the stability and performance of networks based on SRv6. Network failures are inevitable in the real world. Reliability practices can help network engineers quickly identify, isolate, and fix faults, thus minimizing impact on services.

In summary, the necessity of SRv6 reliability practices is evident in several aspects, including improving network stability and performance, enhancing fault handling capabilities, ensuring security, improving compatibility and interoperability, optimizing management and monitoring, and enhancing deployment experience.

SRv6 offers multiple protection mechanisms, with different applications requiring different protection requirements. It is challenging to select the most suitable protection mechanism, or a combination of mechanisms. When multiple protection mechanisms coexist, achieving the desired protection outcome becomes difficult, and there is a lack of effective coordination methods. Some initial work could refer to [I-D.liu-srv6ops-sr-protection].

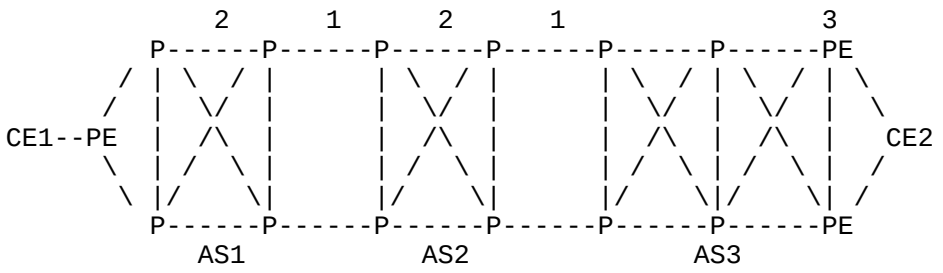

Figure 4: SRv6 Protection Deployment and Networking

Protection includes path protection, intermediate node protection, and service protection, which require different deployment strategies based on their locations. Protection strategies vary for nodes at different locations: for instance, location 1 involves inter-domain link protection, location 2 intra-domain link protection, and location 3 tail node protection.

The selection of path protection strategies also requires consideration of factors such as network architecture, service requirements, resource constraints, and operation expenditure.

DOP-7 SRv6 provides diverse protection mechanisms, but selecting optimal solutions for specific applications remains challenging, especially when coordinating multiple coexisting mechanisms effectively.

7. Challenges of Different Network Types

SRv6 deployment faces various challenges across different network environments, including not only carrier networks but also data centers and campus networks. Due to the diversity of network architectures, traffic patterns, and legacy protocol dependencies, it is difficult to apply a single deployment guideline that meets the unique requirements of each environment.

7.1. Data Center Networks

In large-scale data centers, which commonly adopt Clos (Spine-Leaf) architectures, SRv6 introduces specific challenges:

The architecture relies heavily on Equal-Cost Multi-Path (ECMP) forwarding for traffic balance. SRv6 traffic steering mechanisms,

such as explicit path control or policy-based routing, must be compatible with ECMP to avoid uneven traffic distribution.

- Data centers typically have strict latency and throughput requirements. The SRv6 header overhead may impact overall performance and reduce effective MTU, especially in scenarios with high-throughput workloads.
- Given the massive scale of data centers, there is a strong requirement for automatic SRv6 deployment, streamlined service provisioning, and simplified OAM, to reduce configuration errors and operational burdens.

7.2. Campus Networks

Campus networks present another set of SRv6 deployment challenges, largely driven by the coexistence with traditional technologies and operational constraints:

- Existing campus networks often rely on legacy protocols such as OSPF, VLAN, or other Layer 2 technologies. Integrating SRv6 requires careful consideration of protocol migration strategies or long-term coexistence mechanisms.
- Many campus edge devices and terminals are still IPv4-only, leading to issues during the IPv4-to-IPv6 transition. SRv6 deployment must accommodate dual-stack operation or translation mechanisms.
- SRv6 introduces source routing capabilities, which, if improperly secured, can become new attack surfaces. Enhanced security mechanisms, such as policy validation and access control, are essential.
- Similar to data centers, automated deployment and monitoring tools are needed to reduce manual intervention, simplify SRv6 operations, and improve service assurance in campus environments.

For instance, in the power grid, SRv6 is expected to provide Quality of Service (QoS) guarantees, performance optimization, and other

specialized features to meet stringent reliability and determinism requirements based on the characteristics of power applications.

DOP-8 It is difficult to apply a single deployment guideline to meet the diverse SRv6 requirements of different network types.

Data centers and campus networks each introduce unique constraints, such as ECMP compatibility, performance sensitivity, legacy protocol coordination, IPv4/IPv6 transition, and operational security, all of which must be carefully addressed to enable successful SRv6 adoption.

8. Security Considerations

This document does not introduce additional security concerns. It does not change the security properties of SRv6. For general SRv6 security considerations, see [I-D.ietf-spring-srv6-security].

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017

[RFC4364] E. Rosen and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006

[I-D.liu-srv6ops-sid-address-assignment] Y. Liu and Y. Zhu, "IPv6 Address Assignment for SRv6", Expired, Internet-Draft, draft-liu-srv6ops-sid-address-assignment-01, 25 February 2025, <https://datatracker.ietf.org/doc/html/draft-liu-srv6ops-sid-address-assignment-01>.

[I-D.geng-srv6ops-traffic-steering-to-srv6] G. Geng, Y. Liu, C. Xie and C. Lin, "Best practices for traffic steering to SRv6", Expired, Internet-Draft, draft-geng-srv6ops-traffic-steering-to-srv6-00, 4 March 2024, <https://datatracker.ietf.org/doc/html/draft-geng-srv6ops-traffic-steering-to-srv6-00>.

[I-D.liu-srv6ops-sr-protection] Y. Liu, W. Jiang, C. Lin, X. Geng and Y. Liu, "Operational Guidance for Protection mechanisms in SRv6 Networks", Work in Progress, Internet-Draft, draft-liu-srv6ops-sr-protection-03, 20 March 2025, <https://datatracker.ietf.org/doc/html/draft-liu-srv6ops-sr-protection-03>.

[I-D.ietf-spring-srv6-security] N. Buraglio, T. Mizrahi, T. Tong, L. M. Contreras and F. Gont, "Segment Routing IPv6 Security Considerations", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-security-07, 18 September 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-security-07>.

9.2. Informative References

[I-D.ietf-spring-srv6-mpls-interworking] S. Agrawal, C. Filsfils, D. Voyer, G. Dawra, Z. Li and S. Hegde, "SRv6 and MPLS interworking", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-mpls-interworking-01, 7 July 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-mpls-interworking-01>.

10. Appendix: Possible Missing DOPs for Future Consideration

The following are potential SRv6-specific operational challenges (DOPs) that are currently not covered in the main sections, but may be important for future study and discussion:

DOP-X1: Compression-related operational issues: Header compression (e.g., NEXT-CSID/REPLACE-CSID) introduces challenges in parsing, interoperability, and troubleshooting.
DOP-X2: Operational issues related to network programmability: Programmable SRv6 behaviors increase complexity in validation, consistency, and rollback.
DOP-X3: SID structure and management issues: Ambiguities in locator/function/arguments structure affect policy enforcement, summarization, and domain boundary handling.

DOP-X4: SRv6-specific control plane issues: Includes SRv6 SID advertisement errors, SR Policy lifecycle handling, and multi-vendor feature inconsistencies.
DOP-X5: IPv6 fragmentation handling: SRv6 header size may cause MTU issues, requiring additional operational methods for fragmentation detection and mitigation.
DOP-X6: Route summarization trade-offs: Locator-based summarization may impact visibility, path granularity, and fault localization.
DOP-X7: Common DOPs shared with SR-MPLS (out of scope): Includes controller interaction issues, multiple SBI management, and general SR deployment policies.

Contributors

Daniel Voyer
Bell Canada
Email: Danvoyer@gmail.com

Linjian Song
Alibaba, Inc
Email: linjian.slj@alibaba-inc.com

Satoru Matsushima
SoftBank
Email: satoru.matsushima@g.softbank.co.jp

Chongfeng Xie
China Telecom
Email: xiechf@chinatelecom.cn

Xinxin Yi
China Unicom
Email: yixx3@chinaunicom.cn

Authors' Addresses

Yisong Liu
China Mobile
Email: liuyisong@chinamobile.com

Thomas Graf
Swisscom
Email: Thomas.Graf@swisscom.com

Zoltan Miklos
MTN
Email: Zoltan.Miklos@mtn.com

Luis Contreras
Telefonica
Email: luismiguel.contrerasmurillo@telefonica.com

Nicolai Leymann
Deutsche Telekom
Email: N.Leymann@telekom.de