

SPRING
Internet-Draft
Intended status: Informational
Expires: September 30, 2022

W. Cheng
China Mobile
C. Xie
China Telecom
R. Bonica
Juniper
D. Dukes
Cisco Systems
C. Li
Huawei
P. Shaofu
ZTE
W. Henderickx
Nokia
March 29, 2022

Compressed SRv6 SID List Requirements
draft-ietf-spring-compression-requirement-01

Abstract

This document specifies requirements for solutions to compress SRv6 SID lists.

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 September 30, 2022.

Copyright Notice

Copyright (c) 2022 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Conventions used in this document	4
2.1.	Requirements Language	4
2.2.	Terminology	4
3.	SRv6 SID List Compression Requirements	4
3.1.	Dataplane Efficiency and Performance Requirements	4
3.1.1.	Encapsulation Header Size	5
3.1.2.	Forwarding Efficiency	5
3.1.3.	State Efficiency	6
4.	SRv6 Specific Requirements	6
4.1.	SRv6 Based	6
4.2.	Functional Requirements	7
4.2.1.	SRv6 Functionality	7
4.2.2.	Heterogeneous SID lists	8
4.2.3.	SID list length	8
4.2.4.	SID summarization	8
4.3.	Operational Requirements	9
4.3.1.	Lossless Compression	9
4.3.2.	Preservation of non-routing information	9
4.3.3.	Address Planning	10
4.4.	Scalability Requirements	10
4.4.1.	Adjacency segment scale	10
4.4.2.	Prefix segment scale	11
4.4.3.	Service Scale	11
4.4.4.	Compression Levels	11
5.	Protocol Design Requirements	11
5.1.	SRv6 Base Coexistence	11
6.	Security Requirements	12
6.1.	Security Mechanisms	12
6.2.	SR Domain Protection	12
7.	IANA Considerations	12
8.	Security Considerations	13
9.	Contributors	13
10.	Normative References	13
Appendix A.	Proposed Requirements	14
A.1.	IPv6 Based	14

A.2.	Point to Multipoint	15
A.3.	Parsability	15
Authors' Addresses	15

[1.](#) Introduction

The SPRING working group defined SRv6, with [[RFC8402](#)] describing how the Segment Routing (SR) architecture is instantiated on two data-planes: SR over MPLS (SR-MPLS) and SR over IPv6 (SRv6). SRv6 uses a routing header called the SR Header (SRH) [[RFC8754](#)] and defines SRv6 SID behaviors and a registry for identifying them in [[RFC8986](#)]. SRv6 is a proposed standard and is deployed today.

The SPRING working group has observed that some use cases, such as strict path TE, may require long SRv6 SID lists. There are several proposed methods to reduce the resulting SRv6 encapsulation size by compressing the SID list.

The SPRING working group formed a design team to define requirements for, and analyze proposals to, compress SRv6 SID lists.

It is a goal of the design team to identify solutions to SRv6 SID list compression that are based on the SRv6 standards. As such, this document provides requirements for SRv6 SID list compression solutions that utilize the existing SRv6 data plane and control plane.

It is also a goal of the design team to consider proposals that are not based on the SRv6 data plane and control plane. As such, this document includes requirements to evaluate whether a compression proposal provides all the functionality of SRv6 (section "SRv6 Functionality") in addition to satisfying compression specific requirements.

For each requirement, a description, rationale and metrics are described.

The design team will produce a separate document to analyze the proposals.

This document is a draft; additional requirements are under review, additional requirements will be added, and current requirements may change. [Appendix A](#) contains a subset of requirements without unanimous consensus. Additional requirements without unanimous consensus are not in the appendix.

2. Conventions used in this document

2.1. 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.

2.2. Terminology

SR: Segment Routing

SRH: Segment Routing Header

MPLS: Multiprotocol Label Switching

SR-MPLS: Segment Routing over MPLS data plane

SID: Segment Identifier

SRv6: Segment Routing over IPv6

SRv6 SID List: A list of SRv6 SIDs

Compression proposal: A proposal to compress SRv6 SID lists

SRv6 base: SRv6 as defined in [[RFC8402](#)], [[RFC8754](#)], [[RFC8986](#)]

SID numbering space: may be implemented as

- o a single IGP instance
- o a single IGP level or area
- o two or more autonomous systems that coordinate SID numbering space
- o two or more IGP instances that coordinate SID numbering space

SRv6 Encapsulation Header: The IPv6 header, and any extension headers preceding a payload, used to implement a SRv6 base or compression proposal.

3. SRv6 SID List Compression Requirements

3.1. Dataplane Efficiency and Performance Requirements

3.1.1. Encapsulation Header Size

Description: The compression proposal MUST reduce the size of the SRv6 encapsulation header.

Rationale: A smaller SRv6 encapsulation results in better MTU efficiency.

Metric: Compression is the ratio of the IPv6 encapsulation size of SRv6 as defined in [[RFC8402](#)], [[RFC8754](#)], [[RFC8986](#)] vs the IPv6 encapsulation size of a given proposal. The encapsulation savings of a compression proposal vs the SRv6 base is a useful measurement to compare proposals.

The encapsulation metric (E) records the number of bytes required for a proposal to encapsulate a packet given a specific segment list.

- o $E(\text{proposal, segment list})$.

The encapsulation savings(ES)records the encapsulation savings for a proposal to encapsulate a packet given a specific segment list.

- o $ES(\text{proposal, segment list}) = 1 - E(\text{proposal, segment list})/E(\text{SRv6 base, segment list})$.

3.1.2. Forwarding Efficiency

Description: The compression proposal SHOULD minimize the number of required hardware resources accessed to process a segment.

Rationale: Efficiency in bits on the wire and processing efficiency are both important. Optimizing one at the expense of the other may lead to significant performance impact.

Metric: The data plane efficiency metric (D) records the data plane forwarding efficiency of the proposed solution. Two metrics are used and recorded at each segment endpoint:

- o $D.PRS(\text{segment list})$: number of headers parsed during processing of the segment list, starting from and including the IPv6 header.
- o $D.LKU(\text{segment list})$: number of FIB lookups during processing of the segment list. The type of lookup is also recorded as longest prefix match (LPM) or exact match (EM)

3.1.3. State Efficiency

Description: The compression proposal SHOULD minimize the amount of additional forwarding state stored at a node.

Rationale: Additional state increases the complexity of the control plane and data plane. It can also result in an increase in memory usage.

Metric: The state efficiency metric (S) records the amount of additional forwarding state required by the proposed solution.

- o S(node parameters): the number of additional forwarding states that need to be stored at a node, given a set of node parameters consisting of the number of nodes in the network, number of local interfaces, number of adjacencies. The forwarding state is counted as entries required in a Forwarding Information Base (FIB) at a node.

4. SRv6 Specific Requirements

4.1. SRv6 Based

Description: A solution to compress SRv6 SID Lists SHOULD be based on the SRv6 architecture, control plane and data plane. The compression solution MAY be based on a different data plane and control plane, provided that it derives sufficient benefit.

Rationale: A compression proposal built on existing IETF standards is preferable to creating new standards with equivalent functionality and performance.

Metric: The utilization metric (U) records whether a proposal utilizes the SRv6 specifications.

Utilization is recorded in a table, with a column per proposal and rows consisting of the following metrics:

- o U.[RFC8402](#): utilizes [[RFC8402](#)].
- o U.[RFC8754](#): utilizes [[RFC8754](#)].
- o U.PGM: utilizes [[RFC8986](#)].
- o U.IGP: utilizes [[I-D.ietf-lsr-isis-srv6-extensions](#)].
- o U.BGP: utilizes [[I-D.ietf-bess-srv6-services](#)].
- o U.POL: utilizes [[I-D.ietf-spring-segment-routing-policy](#)].
- o U.BLS: utilizes [[I-D.ietf-idr-bgpls-srv6-ext](#)].
- o U.SVC: utilizes [[I-D.ietf-spring-sr-service-programming](#)].
- o U.OAM: utilizes [[I-D.ietf-6man-spring-srv6-oam](#)].
- o U.ALG: utilizes [[I-D.ietf-lsr-flex-algo](#)].

Each cell contains "yes" for utilizes, or "no" for does not utilize.

4.2. Functional Requirements

4.2.1. SRv6 Functionality

Description: A solution to compress an SRv6 SID list MUST support the functionality of SRv6. This requirement ensures no SRv6 functionality is lost. It is particularly important to understand how a proposal, as evaluated in section "SRv6 Based", provides this functionality.

Rationale: Operators require SRv6 functionality. Evaluating the extent to which a proposal supports SRv6 functionality is important for operators and implementors to understand the impact on network operations.

Metric: The Functionality metric (F) records whether a proposal supports SRv6 functionality and how the functionality is provided.

Functionality is recorded in a table with columns for each proposal and rows consisting of the following metrics:

- o F.SID: Supports SRv6 SID functionality as described in [[RFC8402](#)]
- o F.SCOPE: Supports globally and locally scoped SID functionality as described in [[RFC8402](#)]
- o F.PFX: Supports prefix SID functionality as described in [[RFC8402](#)] and [[RFC8986](#)]
- o F.ADJ: Supports adjacency SID functionality as described in [[RFC8402](#)] and [[RFC8986](#)]
- o F.BIND: Supports binding SID functionality as described in [[RFC8402](#)] and [[RFC8986](#)]
- o F.PEER: Supports BGP peering SID functionality as described in [[RFC8402](#)] and [[RFC8986](#)]
- o F.SVC: Supports L3 and L2 VPN service SID functionality as described in [[RFC8986](#)]
- o F.ALG: Supports flexible algorithms functionality as described in [[I-D.ietf-lsr-flex-algo](#)]
- o F.TILFA: Supports TI-LFA functionality as described in [[I-D.ietf-rtgwg-segment-routing-ti-lfa](#)]
- o F.SEC: Supports securing an SR domain with ingress filtering as functionally defined in [[RFC8754](#)]
- o F.IGP: Supports distributing topological SIDs and behaviors via ISIS as functionally described in [[I-D.ietf-lsr-isis-srv6-extensions](#)]
- o F.BGP: Supports BGP VPNs as functionally described in [[I-D.ietf-bess-srv6-services](#)]

- o F.POL: Supports SR policies and steering traffic over those policies as functionally described in [[I-D.ietf-spring-segment-routing-policy](#)]
- o F.BLS: Supports Link State distribution via BGP as functionally described in [[I-D.ietf-idr-bgpls-srv6-ext](#)]
- o F.SFC: Supports stateless service programming as functionally described in [[I-D.ietf-spring-sr-service-programming](#)]
- o F.PING: Supports pinging a SID to verify the SID is implemented as functionally described in [[I-D.ietf-6man-spring-srv6-oam](#)]

Each cell contains the specification name documenting the functionality.

[4.2.2.](#) Heterogeneous SID lists

Description: The compression proposal SHOULD support a combination of compressed and non-compressed segments in a single path. As an example, a solution may satisfy this requirement without being SRv6 based by using a binding SID to impose an additional SRv6 header (IPv6 header plus optional SRH) with non-compressed SID.

Rationale: Support of SID lists with compressed and non-compressed SIDs reduces encapsulation size when not all SRv6 nodes deploy the compression proposal or 128-bit SIDs are required.

Metric: A compliant compression proposal supports both:

- o classic 128-bit SRv6 SIDs in the IPv6 Destination Address field
- o segment lists (i.e., paths) with both compressed and 128-bit SRv6 SIDs.

[4.2.3.](#) SID list length

Description: The compression proposal MUST be able to represent SR paths that contain up to 16 segments.

Rationale: Strict TE paths require SID list lengths proportional to the diameter of the SR domain.

Metric: The compression proposal must be able to steer a packet through an SR path that contains up to sixteen segments.

[4.2.4.](#) SID summarization

Description: The solution MUST be compatible with segment summarization.

Rationale: Summarization of segments is a key benefit of SRv6 vs SR MPLS. In interdomain deployments, any node can reach any other node via a single prefix segment. Without summarization, border router SIDs must be leaked, and an additional global prefix segment is required for each domain border to be traversed.

Metric: A solution supports summarization when segments can be summarized for advertisement into other IGP domains or levels.

4.3. Operational Requirements

4.3.1. Lossless Compression

Description: A path traversed using a compressed SID list MUST always be the same as the path traversed using the uncompressed SID list if no compression was applied.

Rationale: In SRv6, we can represent a path to meet certain objectives. A compression proposal needs to support the objectives with the same path.

Metric: Information present in the pre-compression segment list MUST also be present in the post-compression SID list.

4.3.2. Preservation of non-routing information

Description: The compression mechanism MUST NOT cause the loss of non-routing information when delivering a packet from the SR ingress node to the egress/penultimate SR node

Rationale: SRv6 ingress nodes encode non-routing information in the IPv6 header chain. This information can be encoded in the following fields:

- o Hop Count
- o DSCP bits
- o ECN bits
- o Flow label
- o HBH Options Extension header
- o Fragment Extension header
- o Authentication Extension header
- o Encrypted Security Payload Extension header
- o Destination Options Extension header

Some of these fields are mutable (e.g., Hop Count) while others are immutable (e.g., Fragment Extension Header).

Some of these fields contain information that is required by every node along a packet's delivery path (e.g., Hop Count). Others contain information that is required only by the packet's ultimate destination (e.g., Fragment Extension Header).

Therefore, the compression mechanism **MUST NOT** prevent this information from being delivered, in an IPv6 header chain, to any node that needs it.

Metric: The SR source node encapsulates its payload (e.g., Ethernet, IP, TCP) in an IPv6 header. The SRv6 header contains both routing and non-routing information. The compression mechanism **MUST NOT** cause the loss of non-routing information when delivering a packet from the SR ingress node to the egress/penultimate SR node.

4.3.3. Address Planning

Description: Network operators require addressing plan flexibility, The compression mechanism **MUST** support flexible IPv6 address planning, it **MUST** support deployment by using GUA from different address blocks.

Rationale: The address planning of the network may vary based on the addressing scheme of the operator, so the solution **MUST** support a flexible addressing scheme. Operators need to deploy the solution based on their own address planning.

Metric: The compression proposal supports locators drawn from different prefixes with the solutions analysis indicating efficiency.

4.4. Scalability Requirements

4.4.1. Adjacency segment scale

Description: The compression proposal **MUST** be capable of representing 65000 adjacency segments per node

Rationale: Typically, network operators deploy networks with tens or hundreds of adjacency segments per node, but some network operators may deploy networks that use more adjacency segments per node.

Metric: A proposal that allows 65000 adjacency segments per node satisfies this requirement.

4.4.2. Prefix segment scale

Description: The compression proposal **MUST** be capable of representing 1 million prefix segments per SID numbering space.

Rationale: Typically, network operators deploy networks with thousands of prefix segments per SID numbering space, but some network operators may deploy networks that use more prefix segments per SID numbering space.

Metric: A proposal that allows 1 million prefix segments per SID numbering space satisfies this requirement.

4.4.3. Service Scale

Description: The compression proposal **MUST** be capable of representing 1 million services per node.

Rationale: Typically, network operators deploy networks with tens to hundreds of thousands of services per node, but some network operators may deploy networks that use more services per node.

Metric: A proposal that allows 1 million services per node satisfies this requirement.

4.4.4. Compression Levels

Description: The compression proposal **SHOULD** be able to support multiple levels of compression.

Rationale: The compression proposal will be deployed in networks of varying size with SID numbering spaces of varying size. Network and service scale can directly impact SID length and the ability of a proposal to compress the SID list.

Metric: A compression proposal that supports relatively better compression for smaller SID numbering spaces and service scale satisfies this requirement.

5. Protocol Design Requirements

5.1. SRv6 Base Coexistence

Description: The compression proposal **MUST** support simultaneous deployment with SRv6 networks.

Rationale: SRv6 is deployed today. A compression proposal that interoperates well with SRv6, as deployed, will reduce the overhead

and simplify operations. For Network operators who would migrate to compressed SRv6 SID lists, the migration is expected to gradually occur over a period of time as they upgrade networks, domains, device families and software instances.

Metric: A compliant compression proposal provides the following

- o Supports simultaneous deployment at a node with current SRv6 SIDs.
- o Supports simultaneous deployment at a node with current SRv6 control plane.
- o Supports simultaneous operation of current SRv6 paths with compressed paths.
- o Supports the behaviors in [[RFC8986](#)].
- o Does not require removal of existing IPv6 address planning.

6. Security Requirements

6.1. Security Mechanisms

Description: The compression solution SHOULD be able to address security issues that it introduces, using existing security mechanisms.

Rationale: It is important to identify security issues and how to address them in any specification.

Metric: A compression solution that does not introduce unresolved security issues meets this requirement.

6.2. SR Domain Protection

Description: A compression solution must not require nodes outside the SR domain to know SID values within the SR domain, and it must provide the ability to block nodes outside an SR domain from accessing SIDS.

Rationale: The unauthorized use of SIDs within the SR domain by nodes outside the domain can disrupt an operators' network.

Metric: A compliant solution describes how access to SIDs within the SR domain is denied to nodes outside the SR domain.

7. IANA Considerations

This document has no requests to IANA.

8. Security Considerations

TBD

9. Contributors

The following individuals contributed to this draft

Sanders Steffann, SJM Steffann Consultancy, sander@steffann.nl

10. Normative References

[I-D.ietf-6man-spring-srv6-oam]

Ali, Z., Filsfils, C., Matsushima, S., Voyer, D., and M. Chen, "Operations, Administration, and Maintenance (OAM) in Segment Routing Networks with IPv6 Data plane (SRv6)", [draft-ietf-6man-spring-srv6-oam-13](#) (work in progress), January 2022.

[I-D.ietf-bess-srv6-services]

Dawra, G., Talaulikar, K., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "SRv6 BGP based Overlay Services", [draft-ietf-bess-srv6-services-15](#) (work in progress), March 2022.

[I-D.ietf-idr-bgpls-srv6-ext]

Dawra, G., Filsfils, C., Talaulikar, K., Chen, M., Bernier, D., and B. Decraene, "BGP Link State Extensions for SRv6", [draft-ietf-idr-bgpls-srv6-ext-09](#) (work in progress), November 2021.

[I-D.ietf-lsr-flex-algo]

Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", [draft-ietf-lsr-flex-algo-18](#) (work in progress), October 2021.

[I-D.ietf-lsr-isis-srv6-extensions]

Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over IPv6 Dataplane", [draft-ietf-lsr-isis-srv6-extensions-18](#) (work in progress), October 2021.

[I-D.ietf-rtgwg-segment-routing-ti-lfa]

Litkowski, S., Bashandy, A., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", [draft-ietf-rtgwg-segment-routing-ti-lfa-08](#) (work in progress), January 2022.

- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-22](#) (work in progress), March 2022.
- [I-D.ietf-spring-sr-service-programming]
Clad, F., Xu, X., Filsfils, C., Bernier, D., Li, C., Decraene, B., Ma, S., Yadlapalli, C., Henderickx, W., and S. Salsano, "Service Programming with Segment Routing", [draft-ietf-spring-sr-service-programming-05](#) (work in progress), September 2021.
- [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>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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>>.

[Appendix A](#). Proposed Requirements

This appendix contains requirements that the design team discussed but could not be agreed upon.

[A.1](#). IPv6 Based

Description: The compression mechanism requires every node along the packet's delivery path to be IPv6-capable. It MUST not require any

node along the packet's forwarding path to support any other forwarding plane (e.g., IPv4, MPLS)

Rational: According to [RFC 8402](#), SRv6 is an instantiation of the SR Architecture over the IPv6 data plane.

Metric: A compliant solution requires every node along the packet's delivery path to be IPv6-capable. It does not require any node along the packet's forwarding path to support any other forwarding plane (e.g., IPv4, MPLS)

[A.2.](#) Point to Multipoint

Description: The compression mechanism SHOULD support point-to-multipoint SR paths.

Rationale: Many VPN services require point-to-multipoint SR paths.

Metric: A compliant proposal can encode a multicast address in the ultimate segment of the segment list.

[A.3.](#) Parsability

Description: The compression mechanism MUST be parsable. That is, the node that consumes the compressed SID list must be able to decode the active and next segment. Parsing information MAY be conveyed in either the forwarding or control plane.

Rationale: Failure to parse the compressed SID list leads to undesired behaviors.

Metric: In the nominal case the producer and consumer of the SID list agree on the active segment and next segment. In foreseeable failure modes it is possible to determine why the producer and consumer don't agree.

Authors' Addresses

Weiqiang Cheng
China Mobile

Email: chengweiqiang@chinamobile.com

Chongfeng Xie
China Telecom

Email: xiechf@chinatelecom.cn

Ron Bonica
Juniper

Email: rbonica@juniper.net

Darren Dukes
Cisco Systems

Email: ddukes@cisco.com

Cheng Li
Huawei

Email: c.l@huawei.com

Peng Shaofu
ZTE

Email: peng.shaofu@zte.com.cn

Wim Henderickx
Nokia

Email: wim.henderickx@nokia.com

