

Workgroup: LSR Working Group

Internet-Draft:

draft-bestbar-lsr-spring-nrp-00

Published: 11 January 2022

Intended Status: Standards Track

Expires: 15 July 2022

Authors: T. Saad V.P. Beeram
 Juniper Networks Juniper Networks
 R. Chen S. Peng B. Wen
 ZTE Corporation ZTE Corporation Comcast
 D. Ceccarelli
 Ericsson

IGP Extensions for SR Network Resource Partition SIDs

Abstract

Segment Routing (SR) defines a set of topological "segments" within an IGP topology to enable steering over a specific SR path. These segments are advertised by the link-state routing protocols (IS-IS and OSPF).

This document describes extensions to the IS-IS and OSPF required to support the signaling of Resource Partition (NRP) segments that operate over SR-MPLS and SRv6 dataplanes. Multiple SR NRP segments can be associated with the same topological element to allow offering of different forwarding treatments (e.g. scheduling and drop policy) associated with each NRP.

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 15 July 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 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](#)
- [2. Requirements Language](#)
- [3. NRP SIDs for SR-MPLS](#)
 - [3.1. IS-IS NRP Prefix-SID Sub-TLV](#)
 - [3.2. IS-IS NRP Adjacency-SID Sub-TLV](#)
 - [3.3. IS-IS NRP per Algorithm Adjacency-SID Sub-TLV](#)
 - [3.4. IS-IS NRP LAN Adjacency-SID Sub-TLV](#)
 - [3.5. IS-IS NRP per Algorithm LAN Adjacency-SID Sub-TLV](#)
- [4. NRP SIDs for SRv6](#)
 - [4.1. SRv6 NRP SID Sub-Sub-TLV](#)
- [5. IANA Considerations](#)
 - [5.1. IS-IS Consideration](#)
 - [5.2. SRv6 IS-IS NRP SID Sub-Sub-TLV](#)
- [6. Security Considerations](#)
- [7. Acknowledgement](#)
- [8. Contributors](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

The Segment Routing (SR) architecture [[RFC8402](#)] defines a set of topological "segments" within an IGP topology as means to enable steering over a specific SR end-to-end path. These segments are advertised by the IGP link-state routing protocols (IS-IS and OSPF). The SR control plane can be applied to both IPv6 and MPLS data planes.

The definition of a network slice for use within the IETF and the characteristics of IETF network slice are specified in [[I-D.ietf-teas-ietf-network-slice-definition](#)]. A framework for reusing IETF VPN and traffic-engineering technologies to realize IETF network slices is discussed in [[I-D.nsd-t-teas-ns-framework](#)].

[[I-D.bestbar-teas-ns-packet](#)] introduces a Slice-Flow Aggregate as the collection of packets (from one or more IETF network slice traffic streams) that match an NRP Policy selection criteria and are offered the same forwarding treatment. The NRP Policy is used to realize an NRP by instantiating specific control and data plane resources on select topological elements in an IP/MPLS network.

[[I-D.bestbar-spring-scalable-ns](#)] describes an approach to extend SR to advertise new SID types called NRP SIDs. Such NRP SIDs are used by a router to define the forwarding action for a packet (next-hop selection), as well as to enforce the specific treatment (scheduling and drop policy) associated with the NRP.

This document defines the IS-IS and OSPF specific encodings for the IGP-Prefix Segment, the IGP-Adjacency Segment, the IGP-LAN-Adjacency Segment that are required to support the signaling of SR NRP SIDs operating over SR-MPLS and SRv6 dataplanes.

When the NRP segments share the same topology (and Algorithm for NRP Prefix-SIDs), the different NRP SIDs of the same topological element share the same forwarding path (i.e., IGP next-hop(s)), but are associated with the specific forwarding treatment (e.g. scheduling and drop policy) of each NRP.

2. 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.

3. NRP SIDs for SR-MPLS

Segment Routing can be directly instantiated on the MPLS data plane through the use of the Segment Routing header instantiated as a stack of MPLS labels defined in [[RFC8402](#)].

3.1. IS-IS NRP Prefix-SID Sub-TLV

[[RFC8667](#)] defines the IS-IS Prefix Segment Identifier sub-TLV (Prefix-SID sub-TLV) that is applicable to SR-MPLS dataplane. The Prefix-SID sub-TLV carries the Segment Routing IGP-Prefix-SID, and is associated with a prefix advertised by a router.

A new IS-IS SR Network Resource Partition Prefix SID (NRP Prefix-SID) sub-TLV is defined to allow a router advertising a prefix to associate multiple NRP Prefix-SIDs to the same prefix. The NRP Prefix-SIDs associated with the same prefix share the same IGP path to the destination prefix within the specific mapped or customized

topology/algorithm but offer the specific QoS treatment associated with the specific NRP.

The NRP ID is carried in the NRP Prefix-SID sub-TLV in order to associate the Prefix-SID with the specific NRP. The NRP Prefix-SID sub-TLV has the following format:

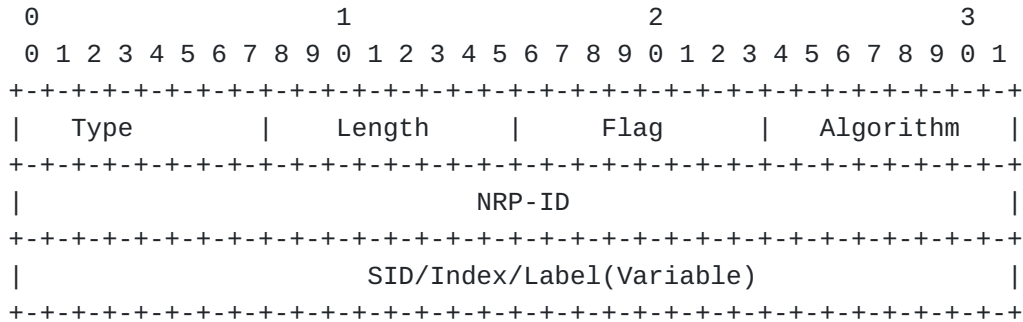


Figure 1: NRP Prefix-SID sub-TLV for SR-MPLS.

where:

Type: TBD1 (Suggested value to be assigned by IANA)

Length: Variable. Depending on the size of the SID.

The "Flags" and "SID/Index/Label" fields are the same as the Prefix-SID sub-TLV [\[RFC8667\]](#).

Algorithm: 1 octet. Associated algorithm. Algorithm values are defined in the IGP Algorithm Type registry

NRP-ID: Identifies a specific NRP within the IGP domain.

This sub-TLV MAY be present in any of the following TLVs:

TLV-135 (Extended IPv4 reachability) defined in [\[RFC5305\]](#).

TLV-235 (Multitopology IPv4 Reachability) defined in [\[RFC5120\]](#).

TLV-236 (IPv6 IP Reachability) defined in [\[RFC5308\]](#).

TLV-237 (Multitopology IPv6 IP Reachability) defined in [\[RFC5120\]](#).

This sub-TLV MAY appear multiple times in each TLV.

3.2. IS-IS NRP Adjacency-SID Sub-TLV

[RFC8667] defines the IS-IS Adjacency Segment Identifier sub-TLV (Adj-SID sub-TLV). The Adj-SID sub-TLV is an optional sub-TLV carrying the Segment Routing IGP Adjacency-SID as defined in [RFC8402].

A new SR Network Resource Partition Adjacency SID (NRP Adj-SID) sub-TLV is defined to allow a router to allocate and advertise multiple NRP Adj-SIDs towards the same IS-IS neighbor (adjacency). The NRP Adj-SIDs allows a router to enforce the specific treatment associated with the NRP on the specific adjacency.

The NRP ID is carried in the NRP Adj-SID sub-TLV to associate it to the specific NRP, and has the following format:

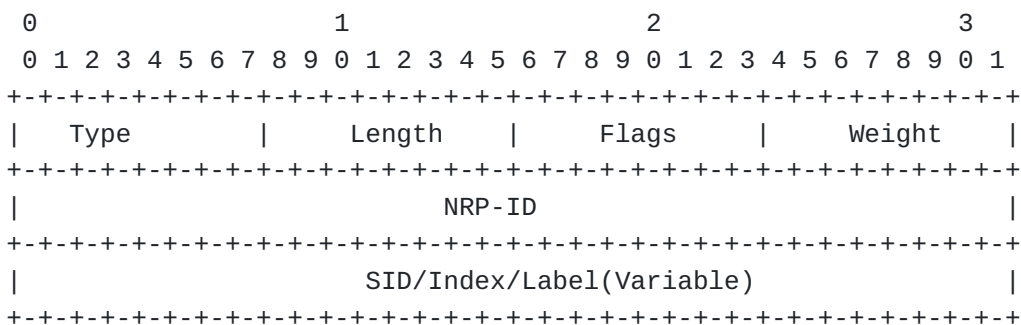


Figure 2: NRP Adj-SID sub-TLV for SR-MPLS.

where:

Type: TBD2 (Suggested value to be assigned by IANA)

Length: Variable. Depending on the size of the SID.

The "Flags", "SID/Index/Label", and "Weight" fields are the same as those defined for the Adj-SID sub-TLV in [RFC8667].

NRP-ID: Identifies a specific NRP within the IGP domain.

This sub-TLV MAY be present in any of the following TLVs:

TLV-22 (Extended IS reachability) [[RFC5305](#)].

TLV-222 (Multitopology IS) [[RFC5120](#)].

TLV-23 (IS Neighbor Attribute) [[RFC5311](#)].

TLV-223 (Multitopology IS Neighbor Attribute) [[RFC5311](#)].

TLV-141 (inter-AS reachability information) [[RFC5316](#)].

Multiple Adj-SID sub-TLVs MAY be associated with a single IS-IS neighbor. This sub-TLV MAY appear multiple times in each TLV.

3.3. IS-IS NRP per Algorithm Adjacency-SID Sub-TLV

[[I-D.ietf-lsr-algorithm-related-adjacency-sid](#)] defines ISIS Adjacency Segment Identifier (Adj-SID) per Algorithm Sub-TLV.

A new per Algorithm SR NRP Adj-SID is defined to allow a router to allocate and advertise multiple NRP Adj-SIDs towards the same adjacency. The per Algorithm NRP Adj-SID allow the router to enforce the specific forwarding treatment associated with the NRP on to packets using that NRP Adj-SID as active segment.

The NRP ID is carried in the NRP per Algorithm Adj-SID sub-TLV to associate it to the specific NRP. The sub-TLV has the following format:




Figure 3: Per Algorithm NRP Adj-SID sub-TLV for SR-MPLS.

where:

Type: TBD3.

Length: 10 or 11 depending on size of the SID.

NRP-ID: Identifies a specific NRP within the IGP domain.

The "Flags", "SID/Index/Label", and "Weight" fields are the same as those defined for the Adj-SID sub-TLV in [RFC8667].

The "Algorithm" field is as defined in [[I-D.ietf-lsr-algorithm-related-adjacency-sid](#)] for the per Algorithm Adj-SID Sub-TLV.

3.4. IS-IS NRP LAN Adjacency-SID Sub-TLV

In LAN subnetworks, [RFC8667] defines the SR-MPLS LAN-Adj-SID sub-TLV for a router to advertise the Adj-SID of each of its neighbors.

A new SR Network Resource Partition LAN Adjacency SID (NRP LAN-Adj-SID) sub-TLV is defined to allow a router to allocate and advertise multiple NRP LAN-Adj-SIDs towards each of its neighbors on the LAN. The NRP LAN-Adj-SIDs allows a router to enforce the specific treatment associated with the specific NRP towards a neighbor.

The NRP ID is carried in the NRP LAN-Adj-SID sub-TLV to associate it to the specific NRP, and it has the following format:

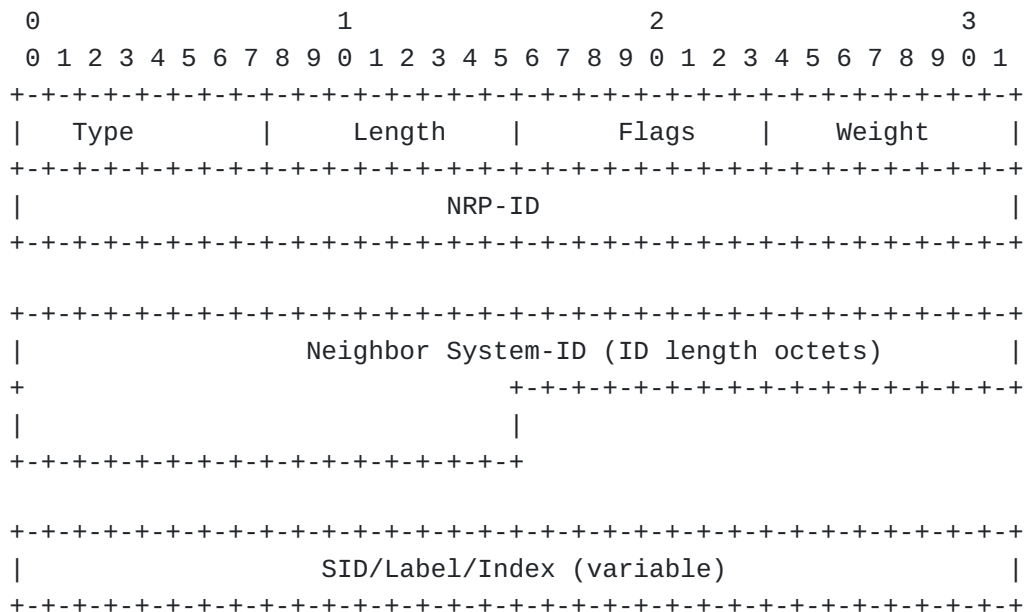


Figure 4: NRP LAN Adj-SID sub-TLV for SR-MPLS.

Type: TBD4 (Suggested value to be assigned by IANA)

Length: Variable. Depending on the size of the SID.

The "Flags" and "SID/Index/Label" fields are the same as the LAN-Adj-SID sub-TLV [[RFC8667](#)].

NRP-ID: Identifies a specific NRP within the IGP domain.

TLV-22 (Extended IS reachability) [[RFC5305](#)].

TLV-23 (IS Neighbor Attribute) [[RFC5311](#)].

Multiple LAN-Adj-SID sub-TLVs MAY be associated with a single IS-IS neighbor. This sub-TLV MAY appear multiple times in each TLV.

ISIS Adjacency Segment Identifier (LAN-Adj-SID) per Algorithm Sub-TLV has the following format:



where:

Type: TBD5.

Length: Variable.

The "Flags", "SID/Index/Label", "Weight", and "Neighbor System-ID" fields are the same as those defined for the LAN-Adj-SID sub-TLV in [[RFC8667](#)].

The "Algorithm" field is as defined in [[I-D.ietf-lsr-algorithm-related-adjacency-sid](#)] for the per Algorithm LAN-Adj-SID Sub-TLV.

Editor Note: the OSPF Sub-TLV sections will be populated in further update.

4. NRP SIDs for SRv6

Segment Routing can be directly instantiated on the IPv6 data plane through the use of the Segment Routing Header defined in [[RFC8754](#)]. SRv6 refers to this SR instantiation on the IPv6 dataplane.

The SRv6 Locator TLV was introduced in [[I-D.ietf-lsr-isis-srv6-extensions](#)] to advertise SRv6 Locators and End SIDs associated with each locator.

4.1. SRv6 NRP SID Sub-Sub-TLV

The SRv6 End SID sub-TLV was introduced in [[I-D.ietf-lsr-isis-srv6-extensions](#)] to advertise SRv6 Segment Identifiers (SID) with Endpoint behaviors which do not require a particular neighbor.

The SRv6 End SID sub-TLV is advertised in the SRv6 Locator TLV, and inherits the topology/algorithm from the parent locator. The SRv6 End SID sub-TLV defined in [[I-D.ietf-lsr-isis-srv6-extensions](#)] carries optional sub-sub-TLVs.

A new SRv6 NRP SID Sub-Sub-TLV is defined to allow a router to assign and advertise an SRv6 End SID that is associated with a specific NRP. The SRv6 SID NRP Sub-Sub-TLV allows routers to infer and enforce the specific treatment associated with the NRP on the selected next-hops along the path to the End SID destination.




Figure 6: SRv6 SID NRP Sub-Sub-TLV format for SRv6.

where:

- Type: TBD6
- Length: 4 octets.
- NRP-ID: Identifies a specific NRP within the IGP domain.

ISIS SRv6 SID NRP Sub-Sub-TLV MUST NOT appear more than once in its parent Sub-TLV. If it appears more than once in its parent Sub- TLV, the parent Sub-TLV MUST be ignored by the receiver.

The new SRv6 SID NRP Sub-Sub-TLV is an optional Sub-Sub-TLV of:

- SRv6 End SID Sub-TLV (Section 7.2 of [[I-D.ietf-lsr-isis-srv6-extensions](#)])
- SRv6 End.X SID Sub-TLV (Section 8.1 of [[I-D.ietf-lsr-isis-srv6-extensions](#)])
- SRv6 LAN End.X SID Sub-TLV (Section 8.2 of [[I-D.ietf-lsr-isis-srv6-extensions](#)])

5. IANA Considerations

This document requests allocation for the following Sub-TLVs types.

5.1. IS-IS Consideration

Table 1 summarizes registrations made in the "Sub-TLVs for TLV 135,235,226 and 237 registry".

Sub-TLV Type	Description	Reference
TBD1	NRP Prefix-SID Sub-TLV	Section 3.1

Table 1

Table 1: Summary of Sub-TLV registrations for TLVs 135,235,226 and 237 (to be assigned by IANA).

Table 2 summarizes registrations made in the "Sub-TLVs for TLV 22, 23, 25, 141, 222, and 223" registry.

Sub-TLV Type	Description	Reference
TBD2	NRP Adj-SID Sub-TLV	Section 3.2
TBD3	NRP LAN-Adj-SID Sub-TLV	Section 3.4
TBD4	NRP Per Algo Adj-SID Sub-TLV	Section 3.3
TBD5	NRP Per Algo LAN-Adj-SID Sub-TLV	Section 3.5

Table 2

Table 2: Summary of Sub-TLV registrations for TLVs 22, 23, 25, 141, 222, and 223 (to be assigned by IANA).

5.2. SRv6 IS-IS NRP SID Sub-Sub-TLV

The below is a request to allocate a new sub-sub-TLV type from the "sub-sub-TLVs for SRv6 End SID and SRv6 End.X SID" registry:

Type: TBD5 (to be assigned by IANA). Reference: [Section 4.1](#)

6. Security Considerations

TBD.

7. Acknowledgement

The authors would like to thank Swamy SRK, and Prabhu Raj Villadathu Karunakaran for their review of this document, and for providing valuable feedback on it.

8. Contributors

The following individuals contributed to this document:

Colby Barth
Juniper Networks
Email: cbarth@juniper.net

Srihari R. Sangli
Juniper Networks
Email: ssangli@juniper.net

Chandra Ramachandran
Juniper Networks
Email: csekar@juniper.net

9. References

9.1. Normative References

[I-D.bestbar-spring-scalable-ns]

Saad, T., Beeram, V. P., Chen, R., Peng, S., Wen, B., and D. Ceccarelli, "Scalable Network Slicing over SR Networks", Work in Progress, Internet-Draft, draft-bestbar-spring-scalable-ns-02, 16 September 2021, <<https://www.ietf.org/archive/id/draft-bestbar-spring-scalable-ns-02.txt>>.

[I-D.bestbar-teas-ns-packet]

Saad, T., Beeram, V. P., Wen, B., Ceccarelli, D., Halpern, J., Peng, S., Chen, R., Liu, X., Contreras, L. M., Rokui, R., and L. Jalil, "Realizing Network Slices in IP/MPLS Networks", Work in Progress, Internet-Draft, draft-bestbar-teas-ns-packet-07, 11 January 2022, <<https://www.ietf.org/archive/id/draft-bestbar-teas-ns-packet-07.txt>>.

[I-D.ietf-lsr-algorithm-related-adjacency-sid]

Peng, S., Chen, R., Talaulikar, K., and P. Psenak, "Algorithm Related IGP-Adjacency SID Advertisement", Work in Progress, Internet-Draft, draft-ietf-lsr-algorithm-related-adjacency-sid-01, 9 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-lsr-algorithm-related-adjacency-sid-01.txt>>.

[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", Work in Progress, Internet-Draft, draft-ietf-lsr-isis-srv6-extensions-18, 20 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-lsr-isis-srv6-extensions-18.txt>>.

[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>>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/

RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.

- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.
- [RFC5311] McPherson, D., Ed., Ginsberg, L., Previdi, S., and M. Shand, "Simplified Extension of Link State PDU (LSP) Space for IS-IS", RFC 5311, DOI 10.17487/RFC5311, February 2009, <<https://www.rfc-editor.org/info/rfc5311>>.
- [RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5316, DOI 10.17487/RFC5316, December 2008, <<https://www.rfc-editor.org/info/rfc5316>>.
- [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>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.
- [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>>.

9.2. Informative References

[I-D.ietf-teas-ietf-network-slice-definition]

Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Definition of IETF Network Slices", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slice-definition-01, 22 February 2021, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slice-definition-01.txt>>.

[I-D.nsdt-teas-ns-framework]

Gray, E. and J. Drake, "Framework for IETF Network Slices", Work in Progress, Internet-Draft, draft-nsdt-teas-ns-framework-05, 2 February 2021, <<https://www.ietf.org/archive/id/draft-nsdt-teas-ns-framework-05.txt>>.

Authors' Addresses

Tarek Saad
Juniper Networks

Email: tsaad@juniper.net

Vishnu Pavan Beeram
Juniper Networks

Email: vbeeram@juniper.net

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn

Shaofu Peng
ZTE Corporation

Email: peng.shaofu@zte.com.cn

Bin Wen
Comcast

Email: Bin_Wen@cable.comcast.com

Daniele Ceccarelli
Ericsson

Email: daniele.ceccarelli@ericsson.com