

IDR
Internet-Draft
Intended status: Standards Track
Expires: 27 July 2024

R. Chen
Sh. Peng
ZTE Corporation
T. Saad
V. Beeram
Juniper Networks
24 January 2024

**BGP-LS Extensions for SR Network Resource Partition SIDs
draft-chen-idr-bgp-ls-transport-slice-06**

Abstract

This document specifies extensions to the BGP Link-state address-family in order to advertise the information of Network Resource Partition SIDs. An external component (e.g., a controller) then can collect Network Resource Partition SIDs in the "northbound" direction. The draft is applicable to both SR-MPLS and SRv6 dataplanes.

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 27 July 2024.

Copyright Notice

Copyright (c) 2024 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](#)
- [2. Requirements Language](#) [3](#)
- [3. NRP SIDs for SR-MPLS](#) [3](#)
 - [3.1. Node Attributes TLV](#) [3](#)
 - [3.2. Link Attribute TLVs](#) [4](#)
 - [3.2.1. NRP-ID list sub-TLV](#) [5](#)
 - [3.2.2. L2 Bundle Member NRP-ID TLV](#) [5](#)
 - [3.2.3. NRP Adjacency-SID TLV](#) [5](#)
 - [3.2.4. NRP per Algorithm Adjacency-SID TLV](#) [6](#)
 - [3.2.5. NRP LAN-Adj-SID TLV](#) [7](#)
 - [3.2.6. NRP per Algorithm LAN-Adj-SID TLV](#) [8](#)
 - [3.3. Prefix Attribute TLVs](#) [9](#)
- [4. NRP SID for SRv6](#) [10](#)
 - [4.1. Router Capabilities for NRPID](#) [10](#)
 - [4.2. SRv6 NRP SID sub-TLV](#) [11](#)
- [5. Acknowledgements](#) [12](#)
- [6. IANA Considerations](#) [12](#)
- [7. Security Considerations](#) [12](#)
- [8. Normative References](#) [12](#)
- Authors' Addresses [14](#)

1. Introduction

Network slicing allows a Service Provider to create independent and logical networks on top of a common or shared physical network infrastructure. [\[I-D.ietf-teas-ietf-network-slices\]](#) provides the definition of a network slice for use within the IETF and discusses the general framework for requesting and operating IETF Network Slices, their characteristics, and the necessary system components and interfaces. It also defines the term "network resource partition(NRP)", which means a set of network resources that are available to carry traffic and meet the SLOs and SLEs.

[\[I-D.ietf-teas-ns-ip-mpls\]](#) proposes the solution to realize network slicing in IP/MPLS networks. 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-lsr-spring-nrp] 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.

In order to satisfy the need for applications that require topological visibility across one area or Autonomous System (AS). This document specifies extensions to the BGP Link-state address-family in order to advertise the information of Network Resource Partition SIDs. An external component (e.g., a controller) then can collect Network Resource Partition SIDs in the "northbound" direction. The draft is applicable to both SR-MPLS and SRv6 dataplanes.

2. 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](#)].

3. NRP SIDs for SR-MPLS

BGP-LS [[RFC7752](#)] defines the link-state NLRI that can be a Node NLRI, a Link NLRI or a Prefix NLRI. The link-state information is mapped to BGP link-state NLRI within the BGP-LS Attribute. In addition, [[I-D.ietf-idr-bgppls-inter-as-topology-ext](#)] defines Stub Link NLRI that is used to describe the inter-as link. This document adds additional BGP-LS Attribute TLVs in order to encode Network Resource Partition(NRP) information. It does not introduce any changes to the encoding of the BGP-LS NLRIs.

3.1. Node Attributes TLV

The NRP-ID TLV is used in order to advertise which Network Resource a router wants to take part in. The NRP-ID sub-TLV is a new TLV of the optional BGP-LS Attribute that is associated with the node NLRI.

The NRP-ID TLV has the following format:




Figure 1

where:

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

Length: variable.

NRPID: Network Resource Partition ID is used to indication the resources on specific link(s)/node(s) that will be traversed by a slice-aggregate.

3.2. Link Attribute TLVs

The following Link Attribute TLVs are are defined:

Type	Description
TBD2	NRP-ID list TLV
TBD3	L2 Bundle Member NRP-ID TLV
TBD4	NRP Adjacency-SID TLV
TBD5	NRP per Algorithm Adjacency-SID TLV
TBD6	NRP LAN-Adj-SID TLV
TBD7	NRP per Algorithm LAN-Adj-SID TLV

Figure 2: xml_happy2

These TLVs should only be added to the BGP-LS Attribute associated with the Link NLRI or Stub Link NLRI.

3.2.1. NRP-ID list sub-TLV

The NRP-ID list sub-TLV has the following format:




Figure 3

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

Length: variable.

NRPID:Network Resource Partition ID is used to indication the resources on specific link(s)/node(s) that will be traversed by a slice-aggregate.

3.2.2. L2 Bundle Member NRP-ID TLV

This TLV is used to advertise NRP-ID for L2 Bundle Member associated with a parent L3 adjacency which is Point-to-Point. The following format is defined for this sub-TLV:

Type: TBD3.

Length: variable.

L2 Bundle Member NRP-ID. There MUST be one NRP-ID for each of the L2 Bundle Members advertised under the preceding L2 Bundle Member Attribute Descriptor. The parent link can be configured to an IGP instance, or as an inter-as link.

3.2.3. NRP Adjacency-SID TLV

This TLV is used to advertise multiple NRP Adj-SID to the controller, Adjacency-SID need to be allocated per NRP-ID. This information is derived from IS-IS NRP Adjacency-SID Sub-TLV (section 3.2 of [\[I-D.bestbar-lsr-spring-nrp\]](#)).

The following format is defined for this sub-TLV:




Figure 4

where:

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

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

Weight: Variable. 1 octet carrying the weight used for load-balancing purposes. The use of weight is described in [section 3.4 of \[RFC8402\]](#).

NRP-ID: Identifies the Network Resource Partition information corresponding to the Adjacency-SID.

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

3.2.4. NRP per Algorithm Adjacency-SID 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 to the controller. The NRP ID is carried in the NRP per Algorithm Adj-SID sub-TLV to associate it to the specific NRP. This information is derived from IS-IS NRP per Algorithm Adjacency-SID Sub-TLV (section 3.3 of [\[I-D.bestbar-lsr-spring-nrp\]](#)).

The following format is defined for this sub-TLV:



Figure 5

where:

Type: TBD5 (Suggested value to be assigned by IANA)

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

Weight: Variable. 1 octet carrying the weight used for load-balancing purposes. The use of weight is described in [section 3.4 of \[RFC8402\]](#).

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

NRP-ID: Identifies the Network Resource Partition information corresponding to the Adjacency-SID.

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

3.2.5. NRP LAN-Adj-SID TLV

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

NRP LAN-Adj-SID TLV is used to advertise multiple NRP LAN-Adj-SIDs, Adjacency-SID need to be allocated per NRPID. This information is derived from the IS-IS NRP LAN Adjacency-SID of IS-IS ([section 3.4 of I-D. \[I-D.bestbar-lsr-spring-nrp\]](#)).

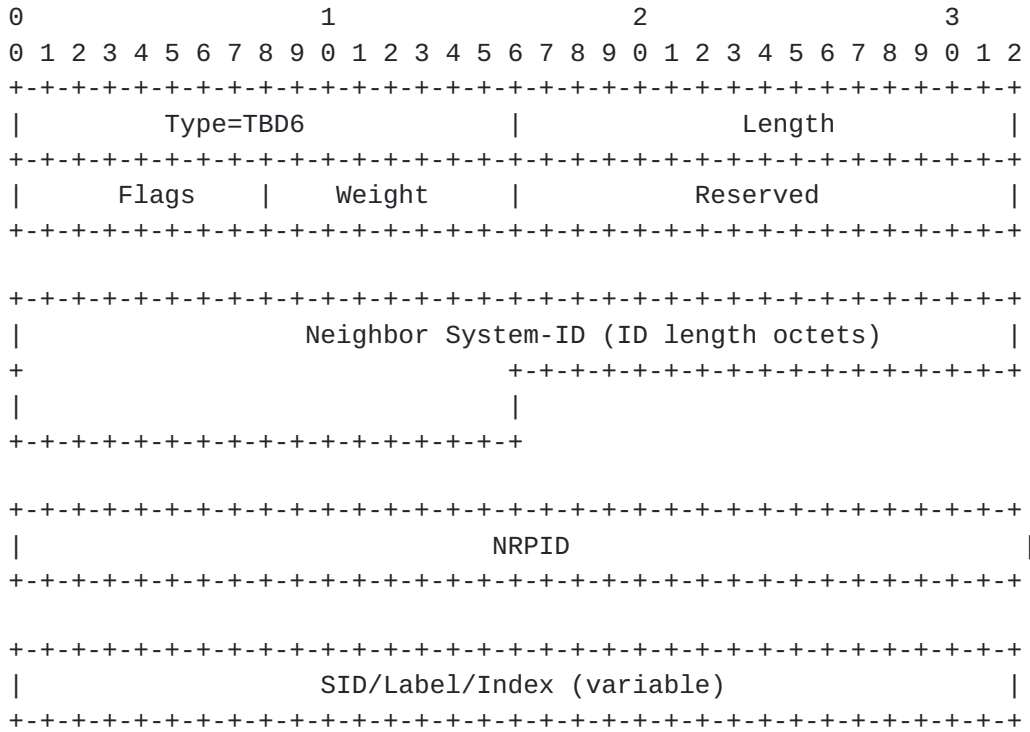


Figure 6

where:

Type:TBD6 (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 Adjacency-SID sub-TLV [RFC8667].

NRP-ID: Identifies the Network Resource Partition information corresponding to the LAN-Adjacency-SID.

3.2.6. NRP per Algorithm LAN-Adj-SID TLV

This information is derived from the IS-IS NRP per Algorithm LAN Adjacency-SID of IS-IS (section 3.5 of I-D. [I-D.bestbar-lsr-spring-nrp]).

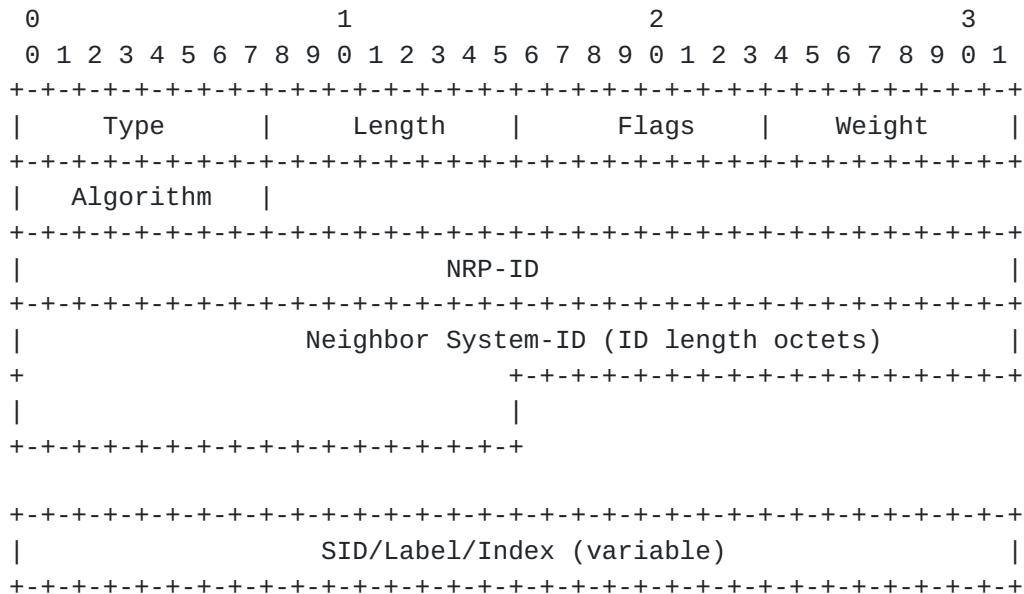


Figure 7

where:

Type:TBD7 (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 Adjacency-SID sub-TLV [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.

NRP-ID: Identifies the Network Resource Partition information corresponding to the LAN-Adjacency-SID.

3.3. Prefix Attribute TLVs

NRP Prefix-SID TLV should only be added to the BGP-LS Attribute associated with the Prefix NLRI describing the prefix of the IGP node.This TLV is used to advertising multiple NRP Prefix-SIDs. This information is derived from IS-IS NRP Prefix-SID Sub-TLV (section 3.1 of [I-D.bestbar-lsr-spring-nrp]).

The NRP Prefix-SID Sub-TLV has the following format:

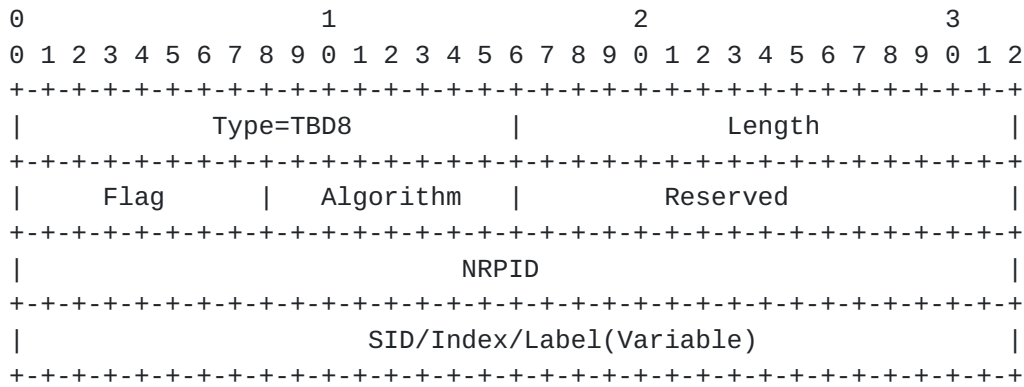


Figure 8

where:

Type:TBD8 (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: Accoridng to section "3.2. SR-Algorithm Sub-TLV" of [RFC8667], two values can be set in this field.

- * 0: Shortest Path First (SPF) algorithm based on link metric.
- * 1: Strict Shortest Path First (SPF) algorithm based on link metric.

NRP-ID: Identifies the NRPID information corresponding to the Prefix-SID.

4. NRP SID for SRv6

SRv6 attributes with an IPv6 prefix are advertised using the new BGP-LS Attribute TLVs defined in this section and associated with the BGP-LS Prefix NLRI.

4.1. Router Capabilities for NRPID

This BGP-LS Attribute TLV is used to announce which NRP-ID a router wants to take part in.

The Router Capabilities for NRPID has the following format:

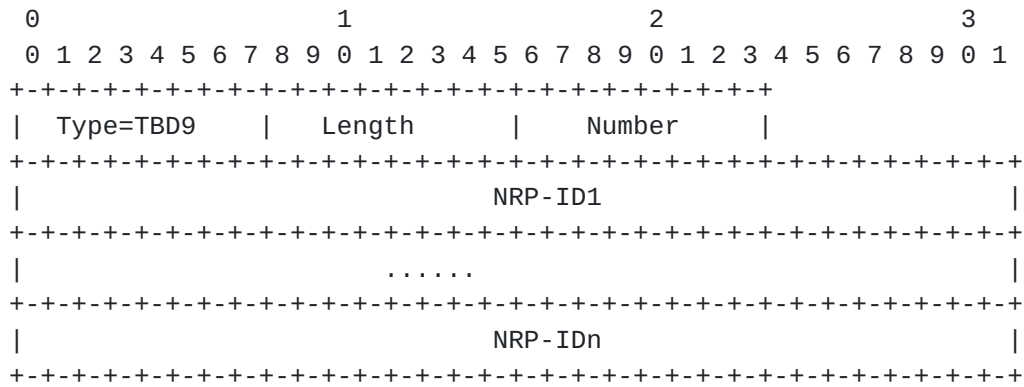


Figure 9

where:

Type: TBD9.

Length: variable.

Number: Number of slice-aggregate which share the same topology.

NRP-ID: Network Resource Partition ID is used to indication the resources on specific link(s)/node(s) that will be traversed by a slice-aggregate.

4.2. SRv6 NRP SID sub-TLV

The SRv6 Locator TLV was introduced in [RFC9514] to advertise SRv6 Locators and additional attributes for the given SRv6 Locator. A new SRv6 NRP SID sub-TLV under the SRv6 Locator TLV is defined to advertise an Locator that is associated with a specific NRP-ID.

The SRv6 NRP SID sub-TLV has the following format:




Figure 10

where:

Type: TBD10

Length: 4 octets.

NRPID: Identifies the NRPID information corresponding to an Locator.

The new NRPID SID sub-TLV is an optional Sub-TLV of:

SRv6 End.X SID TLV ([Section 4.1 of \[RFC9514\]](#)).

SRv6 LAN End.X SID TLV ([Section 4.2 of \[RFC9514\]](#)).

5. Acknowledgements

TBD.

6. IANA Considerations

TBD.

7. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the "Security Considerations" section of [\[RFC4271\]](#) for a discussion of BGP security. Also, refer to [\[RFC4272\]](#) and [\[RFC6952\]](#) for analyses of security issues for BGP. Security considerations for acquiring and distributing BGP-LS information are discussed in [\[RFC7752\]](#).

8. Normative References

[I-D.bestbar-lsr-spring-nrp]

Saad, T., Beeram, V. P., Chen, R., Peng, S., Wen, B., and D. Ceccarelli, "IGP Extensions for SR Network Resource Partition SIDs", Work in Progress, Internet-Draft, [draft-bestbar-lsr-spring-nrp-01](#), 11 July 2022, <<https://datatracker.ietf.org/doc/html/draft-bestbar-lsr-spring-nrp-01>>.

[I-D.ietf-idr-bgpls-inter-as-topology-ext]

Wang, A., Chen, H., Talaulikar, K., and S. Zhuang, "BGP-LS Extension for Inter-AS Topology Retrieval", Work in Progress, Internet-Draft, [draft-ietf-idr-bgpls-inter-as-topology-ext-13](#), 3 April 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgpls-inter-as-topology-ext-13>>.

[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-06](https://datatracker.ietf.org/doc/html/draft-ietf-lsr-algorithm-related-adjacency-sid-06), 5 December 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-algorithm-related-adjacency-sid-06>>.

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

Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", Work in Progress, Internet-Draft, [draft-ietf-lsr-flex-algo-26](https://datatracker.ietf.org/doc/html/draft-ietf-lsr-flex-algo-26), 17 October 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-flex-algo-26>>.

[I-D.ietf-teas-ietf-network-slices]

Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", Work in Progress, Internet-Draft, [draft-ietf-teas-ietf-network-slices-25](https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slices-25), 14 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slices-25>>.

[I-D.ietf-teas-ns-ip-mpls]

Saad, T., Beeram, V. P., Dong, J., Wen, B., Ceccarelli, D., Halpern, J. M., 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-ietf-teas-ns-ip-mpls-03](https://datatracker.ietf.org/doc/html/draft-ietf-teas-ns-ip-mpls-03), 26 November 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ns-ip-mpls-03>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](https://datatracker.ietf.org/doc/html/bcp-14), [RFC 2119](https://datatracker.ietf.org/doc/html/rfc2119), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](https://datatracker.ietf.org/doc/html/rfc4271), DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", [RFC 4272](https://datatracker.ietf.org/doc/html/rfc4272), DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.

- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", [RFC 6952](#), DOI 10.17487/RFC6952, May 2013, <<https://www.rfc-editor.org/info/rfc6952>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", [RFC 7752](#), DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [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>>.
- [RFC9514] Dawra, G., Filsfils, C., Talaulikar, K., Ed., Chen, M., Bernier, D., and B. Decraene, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing over IPv6 (SRv6)", [RFC 9514](#), DOI 10.17487/RFC9514, December 2023, <<https://www.rfc-editor.org/info/rfc9514>>.

Authors' Addresses

Ran Chen
ZTE Corporation
Nanjing
China
Email: chen.ran@zte.com.cn

Shaofu Peng
ZTE Corporation
Nanjing
China
Email: peng.shaofu@zte.com.cn

Tarek Saad
Juniper Networks
Email: tsaad@juniper.net

Vishnu Pavan Beeram
Juniper Networks
Email: vbeeram@juniper.net