

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2019

H. Chen
Z. Li
S. Zhuang
Huawei
March 8, 2019

BGP Extensions for SRv6 SIDs Allocation
draft-chen-idr-bgp-srv6-sid-allocation-00

Abstract

This document describes extensions to the BGP-LS for IDs allocation. The IDs are SIDs for segment routing for IPv6 (SRv6). They are distributed to their domains if needed.

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](#)].

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 9, 2019.

Copyright Notice

Copyright (c) 2019 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	2
2.	Terminology	2
3.	Protocol Extensions	3
3.1.	Node NLRI for IDs Allocation	3
3.2.	Link NLRI for IDs Allocation	6
4.	IANA Considerations	9
5.	Security Considerations	9
6.	Acknowledgements	9
7.	References	9
7.1.	Normative References	9
7.2.	Informative References	11
	Authors' Addresses	11

[1.](#) Introduction

In a network with a central controller, the controller has the link state information of the network, including traffic engineering information. In addition, the controller allocates and manages the resources of the network in general. It is natural and beneficial for the controller to allocate and manage IDs as a kind of network resources.

This document proposes extensions to the BGP-LS as a controller for allocating Segment Identifiers (SIDs) for segment routing for IPv6 (SRv6). If needed, some SIDs will be distributed into their network domains.

[2.](#) Terminology

The following terminology is used in this document.

SR: Segment Routing.


SRv6: SR for IPv6

SR-Path: Segment Routing Path.

TED: Traffic Engineering Database.

This section defines a new Protocol-ID, called IDs-Allocation (IDA), in the Protocol-ID field of Link State NLRI. The use of the new Protocol-ID allows separation and differentiation between the NLRIs carrying IDs Allocation information from the NLRIs carrying IGP link-state information defined in [RFC7752].

The Node NLRI with the new Protocol-ID is used for allocating the IDs associated with a node. It has the same format (refer to the Figure below) as that defined in [RFC7752] and may contain the descriptor and attributes defined in [RFC7752].



```

~                               Local Node Descriptor                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               Node IDs Allocation                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Where:

Protocol-ID: set to IDs-Allocation, a new Protocol-ID.

Local Node Descriptor: defined in [\[RFC7752\]](#), can be reused.

Node IDs Allocation: TLVs or sub-TLVs for IDs associated with the node given by the Local Node Descriptor.

Node IDs Allocation field may contain some of the followings:


SR-Capabilities TLV (1034): It contains the Segment Routing Global Base (SRGB) range(s) allocated for the node.

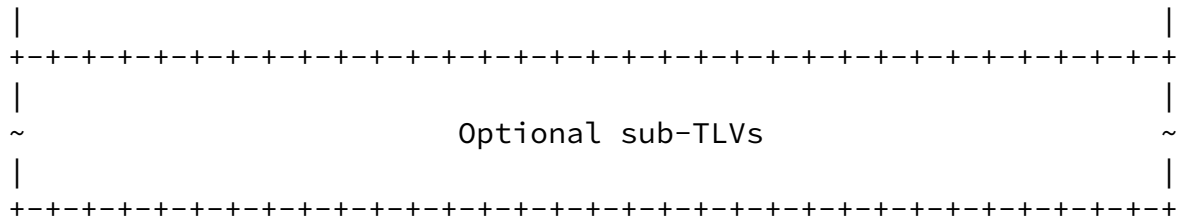
SR Local Block TLV (1036): The SR Local Block (SRLB) TLV contains the range(s) of SIDs/labels allocated to the node for local SIDs.

SRv6 SID Node TLV (TBD1): A new TLV, called SRv6 Node SID TLV, contains an SRv6 SID and related information.

SRv6 Locator TLV (TBD2): A new TLV, called SRv6 Locator TLV, contains an SRv6 locator and related information.

The format of SRv6 SID Node TLV is illustrated below.





SRv6 Node SID TLV

Type: TBD1 for SRv6 Node SID TLV is to be assigned by IANA.

Length: Variable.

Flags: 1 octet. No flags are defined now.

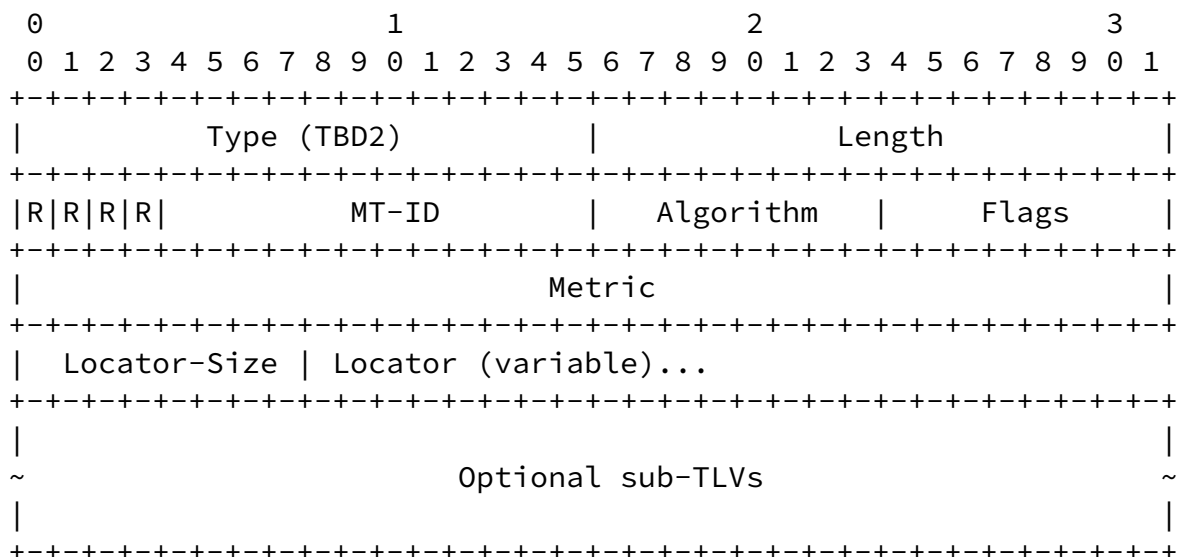
SRv6 Endpoint Function: 2 octets. The function associated with SRv6 SID.

SRv6 Identifier: 16 octets. IPv6 address representing SRv6 SID.

Reserved: MUST be set to 0 while sending and ignored on receipt.

SRv6 node SID inherits the topology and algorithm from its locator.

The format of SRv6 locator TLV is illustrated below.



SRv6 Locator TLV

Type: TBD2 for SRv6 Locator TLV is to be assigned by IANA.

Length: Variable.

MT-ID: Multitopology Identifier as defined in [[RFC5120](#)].

Algorithm: 1 octet. Associated algorithm.

Flags: 1 octet. As described in [\[I-D.bashandy-isis-srv6-extensions\]](#).

Metric: 4 octets. As described in [\[RFC5305\]](#).

Locator-Size: 1 octet. Number of bits in the Locator field (1 to 128).

Locator: 1 to 16 octets. SRv6 Locator encoded in the minimum number of octets for the given Locator-Size.

Reserved: MUST be set to 0 while sending and ignored on receipt.

Chen, et al.

Expires September 9, 2019

[Page 5]


Internet-Draft

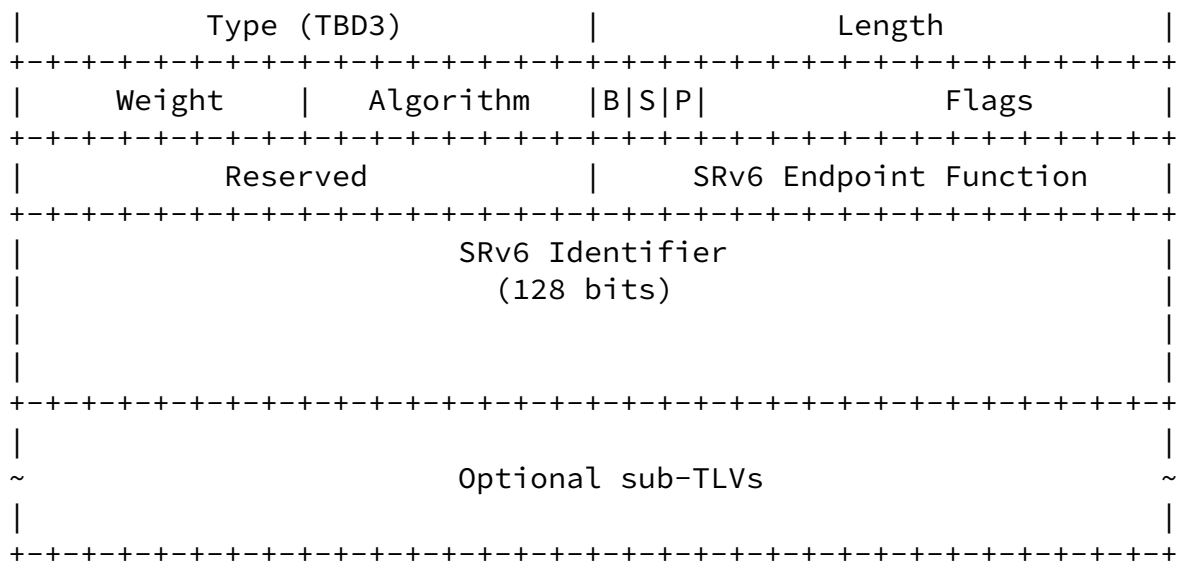
BGP for SRv6 SIDs

March 2019

3.2. Link NLRI for IDs Allocation

The Link NLRI with the new Protocol-ID is used for allocating the IDs associated with a link. It has the same format (refer to the Figure below) as that defined in in [\[RFC7752\]](#) and may contain the descriptors and attributes defined in [\[RFC7752\]](#).





SRv6 Adj-SID TLV

Type: TBD3 for SRv6 Adj-SID TLV is to be assigned by IANA.

Length: Variable.

Weight: 1 octet. The value represents the weight of the SID for the purpose of load balancing.

Algorithm: 1 octet. Associated algorithm.

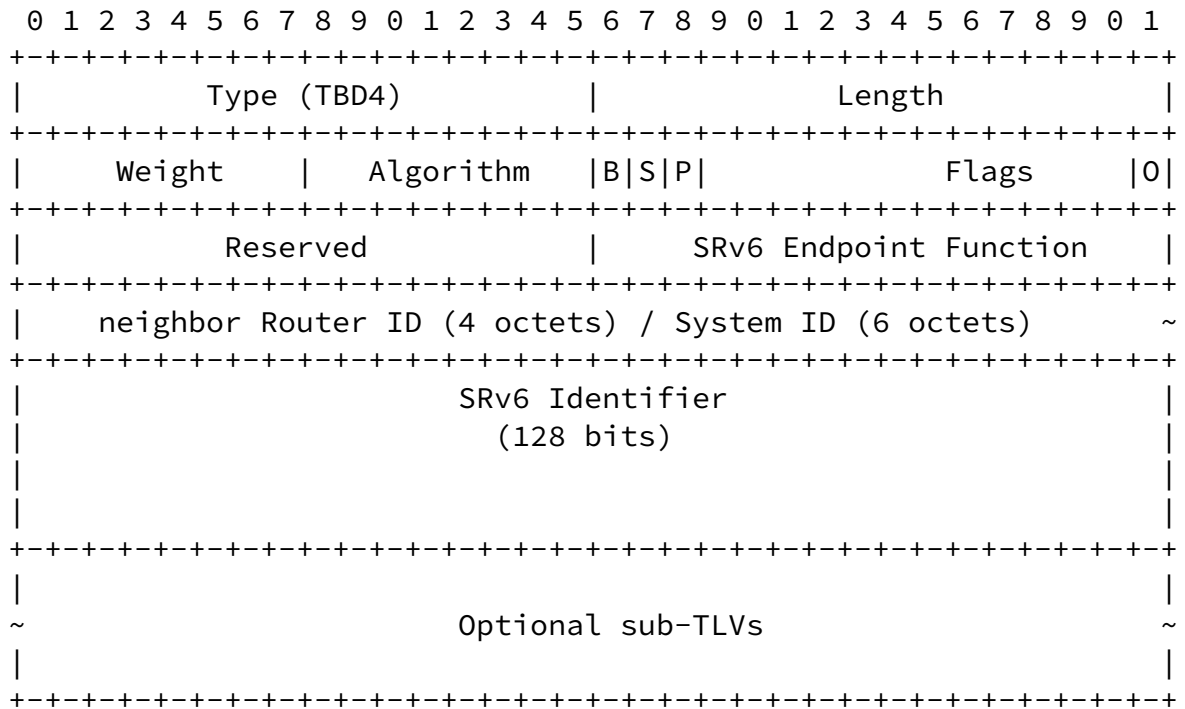
Flags: 2 octets. Three flags are defined in [\[I-D.bashandy-isis-srv6-extensions\]](#).

SRv6 Endpoint Function: 2 octets. The function associated with SRv6 SID.

SRv6 Identifier: 16 octets. IPv6 address representing SRv6 SID.

Reserved: MUST be set to 0 while sending and ignored on receipt.

The format of an SRv6 LAN Adj-SID TLV is illustrated below.



SRv6 LAN Adj-SID TLV

Type: TBD4 for SRv6 LAN Adj-SID TLV is to be assigned by IANA.

Length: Variable.

Weight: 1 octet. The value represents the weight of the SID for the purpose of load balancing.

Algorithm: 1 octet. Associated algorithm.

Flags: 2 octets. Three flags B, S and P are defined in [\[I-D.bashandy-isis-srv6-extensions\]](#). Flag 0 set to 1 indicating OSPF neighbor Router ID of 4 octets, set to 0 indicating IS-IS neighbor System ID of 6 octets.

SRv6 Endpoint Function: 2 octets. The function associated with SRv6 SID.

SRv6 Identifier: 16 octets. IPv6 address representing SRv6 SID.

Reserved: MUST be set to 0 while sending and ignored on receipt.

4. IANA Considerations

This document requests assigning a code-point from the registry "BGP-LS Protocol-IDs" as follows:

Protocol-ID	Description	Reference
TBD	IDs Allocation	Section 3

This document requests assigning a code-point from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" as follows:

TLV Code Point	Description	Reference
TBD1	SRv6 Node SID	Section 3
TBD2	SRv6 Adj-SID	Section 3
TBD3	SRv6 LAN Adj-SID	Section 3

5. Security Considerations

Protocol extensions defined in this document do not affect the BGP security other than those as discussed in the Security Considerations section of [\[RFC7752\]](#).

6. Acknowledgements

The authors would like to thank Nan Wu, and others for their valuable suggestions and comments on this draft.

7. References

7.1. Normative References

[I-D.bashandy-isis-srv6-extensions]
Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Routing over IPv6 Dataplane", [draft-bashandy-isis-srv6-extensions-05](#) (work in progress), March 2019.

Internet-Draft

BGP for SRv6 SIDs

March 2019

[I-D.ietf-idr-flowspec-path-redirect]

Velde, G., Patel, K., and Z. Li, "Flowspec Indirection-id Redirect", [draft-ietf-idr-flowspec-path-redirect-07](#) (work in progress), December 2018.

[I-D.ietf-isis-segment-routing-extensions]

Previdi, S., Ginsberg, L., Filshil, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", [draft-ietf-isis-segment-routing-extensions-22](#) (work in progress), December 2018.

[I-D.ietf-rtgwg-bgp-routing-large-dc]

Lapukhov, P., Premji, A., and J. Mitchell, "Use of BGP for routing in large-scale data centers", [draft-ietf-rtgwg-bgp-routing-large-dc-11](#) (work in progress), June 2016.

[I-D.ietf-spring-segment-routing]

Filshil, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-15](#) (work in progress), January 2018.

[I-D.ietf-spring-segment-routing-ldp-interop]

Bashandy, A., Filshil, C., Previdi, S., Decraene, B., and S. Litkowski, "Segment Routing interworking with LDP", [draft-ietf-spring-segment-routing-ldp-interop-15](#) (work in progress), September 2018.

[I-D.li-ospf-ospfv3-srv6-extensions]

Li, Z., Hu, Z., Cheng, D., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", [draft-li-ospf-ospfv3-srv6-extensions-03](#) (work in progress), March 2019.

[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>>.

Chen, et al.

Expires September 9, 2019

[Page 10]

Internet-Draft

BGP for SRv6 SIDs

March 2019

- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](#), DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.
- [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>>.

[7.2](#). Informative References

- [I-D.gredler-idr-bgp-ls-segment-routing-extension]
Gredler, H., Ray, S., Previdi, S., Filsfils, C., Chen, M., and J. Tantsura, "BGP Link-State extensions for Segment Routing", [draft-gredler-idr-bgp-ls-segment-routing-extension-02](#) (work in progress), October 2014.
- [I-D.ietf-idr-bgpls-segment-routing-epe]
Previdi, S., Talaulikar, K., Filsfils, C., Patel, K., Ray, S., and J. Dong, "BGP-LS extensions for Segment Routing BGP Egress Peer Engineering", [draft-ietf-idr-bgpls-segment-routing-epe-17](#) (work in progress), October 2018.

Authors' Addresses

Huaimo Chen
Huawei
Boston, MA
USA

Email: Huaimo.chen@huawei.com

Zhenbin Li
Huawei
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Chen, et al.

Expires September 9, 2019

[Page 11]

Internet-Draft

BGP for SRv6 SIDs

March 2019

Shunwan Zhuang
Huawei
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: zhuangshunwan@huawei.com

