

IDR Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 19 April 2026

C. Xie  
C. Li  
China Telecom  
J. Dong  
Z. Li  
Huawei Technologies  
16 October 2025

Applicability of Border Gateway Protocol - Link State (BGP-LS) with  
Multi-Topology (MT) for Segment Routing based Network Resource  
Partitions (NRPs)  
draft-ietf-idr-bgpls-sr-vtn-mt-14

## Abstract

When Segment Routing (SR) is used for building Network Resource Partitions (NRPs), each NRP can be allocated with a group of Segment Identifiers (SIDs) to identify the topology and resource attributes of network segments in the NRP. This document describes how BGP-Link State (BGP-LS) with Multi-Topology (MT) can be used to distribute the information of SR-based NRPs to a network controller in a specific context where each NRP is associated with a separate logical topology identified by a Multi-Topology ID (MT-ID). This document sets out the targeted scenarios for the approach suggested, and presents the scalability limitations that arise.

## 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 19 April 2026.

## Copyright Notice

Copyright (c) 2025 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. Advertisement of Topology Information . . . . .	4
2.1. Intra-domain Topology Advertisement . . . . .	4
2.2. Inter-Domain Topology Advertisement . . . . .	4
3. Advertisement of Resource related TE Attribute . . . . .	6
4. Scalability Considerations . . . . .	6
5. Security Considerations . . . . .	7
6. IANA Considerations . . . . .	7
7. Acknowledgments . . . . .	7
8. References . . . . .	7
8.1. Normative References . . . . .	7
8.2. Informative References . . . . .	9
Authors' Addresses . . . . .	9

## 1. Introduction

[RFC9543] discusses the general framework, components, and interfaces for requesting and operating network slices using IETF technologies. [RFC9543] also introduces the concept of the Network Resource Partition (NRP), which is defined as a subset of the buffer/queuing/scheduling resources and associated policies on each of a connected set of links in an underlay network. An NRP can be associated with a logical network topology to select or specify the set of links and nodes involved. [RFC9732] specifies a framework of NRP-based enhanced VPNs and describes the candidate component technologies in different network planes and network layers. An NRP could be used as the underlay to meet the requirements of one or a group of network slices or enhanced VPN services. The mechanism of enforcing NRP resource allocation and the mechanism of mapping one or group of enhanced VPN services to a specific NRP are outside the scope of this document. Similarly, classification and means to bind packets to NRPs are out of scope this document.

[I-D.ietf-spring-resource-aware-segments] introduces resource awareness to Segment Routing (SR) [RFC8402]. As described in [I-D.ietf-spring-sr-for-enhanced-vpn], a group of resource-aware SIDs can be used to build SR-based NRPs with the required network topology

and network resource attributes. The group of resource-aware SR SIDs together with the associated topology and resource attributes of an NRP need to be distributed in the network using IGPs. BGP-Link State (BGP-LS) [RFC9552] can then be used to advertise the SR SIDs and the resource related Traffic Engineering (TE) attributes (e.g., link bandwidth) of NRPs in each IGP area or AS to a network controller.

As specified in [RFC9543], some NRP realizations may build NRPs with dedicated topologies, while other realizations may use a shared topology for multiple NRPs. The exact NRPs characterization, the number of NRPs, and their binding to a topology are deployment-specific.

In some network scenarios, the required number of NRPs may be small, each NRP can be associated with a separate logical topology, i.e., there is 1:1 mapping between an NRP and a Multi-Topology (MT) ID, and the set of dedicated or shared network resources of the NRP can be considered to be associated with the logical topology. [I-D.ietf-lsr-isis-sr-vtn-mt] describes how IS-IS Multi-Topology (MT) [RFC5120] can be used to advertise an independent topology and the associated SR SIDs, together with the topology-specific resource related TE attributes in the network.

In some network scenarios, for instance, an operator's network consists of multiple network parts, such as metro area networks, backbone networks, or data center networks, each part being a different AS. NRPs can be enabled independently in each network part. Specifically, it is not required to enable multiple NRPs in all network parts, and the number of NRPs is local to each domain. However, there might be a need to stitch NRPs that span multiple ASes, typically under the same network administration. NRP stitching is likely to require classification, (re)marking, admission control, etc. at ingress nodes of network domains. These considerations are out of scope of this document.

This document describes how BGP-LS with MT can be used distribute information of the logical topology, the associated SIDs, and the topology-specific resource information to a network controller for intra-domain and inter-domain SR-based NRPs. The limitations and the targeted scenario of this approach are described in "Scalability Considerations" (Section 4).

## 2. Advertisement of Topology Information

This section describes the corresponding BGP-LS mechanism to distribute both the intra-domain and inter-domain topology information and the associated SR SIDs. For the inter-domain case, the involved network domains should be under a common administration, or they belong to the same trusted domain as specified in Section 8 of [RFC8402].

### 2.1. Intra-domain Topology Advertisement

Section 5.2.2.1 of [RFC9552] defines the Multi-Topology Identifier (MT-ID) TLV, which can contain one or more Multi-Topology Identifiers for a link, node, or prefix. The MT-ID TLV may be included as a Link Descriptor, as a Prefix Descriptor, or in the BGP-LS Attribute of a Node Network Layer Reachability Information (NLRI). The detailed rules of the usage of MT-ID TLV in BGP-LS is also specified in Section 5.2.2.1 of [RFC9552].

When Multi-Topology is used with the SR-MPLS data plane [RFC8660], topology-specific Prefix-SIDs and topology-specific Adjacency Segment Identifiers (Adj-SIDs) can be carried in the BGP-LS Attribute associated with the Prefix NLRI and Link NLRI respectively (Section 2 of [RFC9085]), the MT-ID TLV carried in the Prefix Descriptor or Link Descriptor [RFC9552] can be used to identify the corresponding topology of the SIDs.

When Multi-Topology is used with the SRv6 data plane [RFC8754], the SRv6 Locator TLV is carried in the BGP-LS Attribute associated with the Prefix NLRI, the MT-ID TLV can be carried as a Prefix Descriptor to identify the corresponding topology of the SRv6 Locator (Section 6 of [RFC9514]). The SRv6 End.X SIDs are carried in the BGP-LS Attribute associated with the Link NLRI, the MT-ID TLV can be carried in the Link Descriptor to identify the corresponding topology of the End.X SIDs. The SRv6 SID NLRI is defined to advertise other types of SRv6 SIDs, in which the SRv6 SID descriptors can include the MT-ID TLV so as to advertise topology-specific SRv6 SIDs.

### 2.2. Inter-Domain Topology Advertisement

[RFC9086] defines the BGP-LS extensions for BGP EPE with SR-MPLS. The BGP-LS extensions for Egress Peer Engineering with SRv6 are specified in [RFC9514]. These extensions could be used by a network controller for the collection of inter-domain topology and SR SID information, which can be used for the computation and instantiation of inter-AS SR-TE paths.

For the case of inter-domain, the inter-domain connectivity and the BGP peering SR SIDs associated with each logical topology on the inter-domain links need to be advertised. This section describes the applicability of multi-topology for the advertisement of inter-domain topology and the associated SR SIDs using BGP-LS. It does not introduce multi-topology into the operation of BGP sessions on the inter-domain links.

In this document, consistent allocation of MT-ID means that the same MT-ID is used in multiple domains for the concatenation of an inter-domain logical topology. When a MT-ID is consistently allocated in multiple domains, the MT-ID can be carried in the link NLRI of the inter-domain links for the advertisement of inter-domain logical topology and the topology-specific BGP peering SIDs. This can be achieved with the combination of existing mechanisms as defined in [RFC9552][RFC9086] and [RFC9514].

Depending on the different inter-domain scenarios, the approaches for the inter-domain topology advertisement can be one or multiple of the following:

- \* One External BGP (EBGP) session between two ASes can be established over multiple underlying links. In this case, different underlying links may be assigned to different inter-domain logical topologies. In another similar case, the EBGP session is established over a single physical link, while the network resource (e.g., bandwidth) on this link is partitioned, each of which is instantiated as a logical sub-interface, and each logical link can be associated with a separate logical topology. Different BGP Peer-Adj-SIDs or SRv6 End.X SIDs need to be allocated and provisioned to each underlying physical or logical link. The association between the underlying physical or logical links and the corresponding MT-ID, together with the BGP Peer-Adj-SIDs or SRv6 End.X SID need to be advertised by the ASBR to a network controller.
- \* For inter-domain connection(s) between two ASes, multiple EBGP sessions can be established between different sets of peering ASBRs: It is possible that some of these BGP peers are only used for one inter-domain logical topology, while some other BGP peers are used for another inter-domain logical topology. In this case, different BGP Peer Node SIDs can be allocated and provisioned to steer traffic to a specific peer within an inter-domain logical topology. The association between the link of the BGP peering session and the corresponding MT-ID, together with the BGP Peer Node SIDs need to be advertised by the ASBR to a network controller.

- \* At the level inter-AS topology, different inter-domain logical topologies may have different inter-AS connectivity. Different BGP Peer Set SIDs may be allocated and provisioned to represent a group of BGP peers which can be used for load-balancing within each inter-domain logical topology. The BGP Peer Set SIDs may be advertised in the BGP-LS attributes of the link NLRI which carries the corresponding MT-ID.

In network scenarios where consistent allocation of MT-ID among multiple domains for an inter-domain logical topology can not be achieved, the MT-IDs advertised by the two peering ASBRs to a network controller for the same inter-domain link in a logical topology could be different. The ASBRs just need to distribute the inter-domain link information with MT-ID to the controllers, it is the controller's job to provide some mapping mechanism to match the different MT-IDs of an inter-domain link in two directions (e.g., for one inter-domain link, MT-ID A in domain X will be matched with MT-ID B in domain Y), and concatenate the inter-domain logical topology. The detailed mechanism is out of the scope of this document.

### 3. Advertisement of Resource related TE Attribute

The information of the network resources attributes of a link associated with a specific topology can be specified by carrying the corresponding TE Link attribute TLVs in BGP-LS Attribute [RFC9552], with the associated MT-ID carried in the corresponding Link NLRI.

For example, a subset of the bandwidth resource on a link for a specific logical topology can be advertised by carrying the Maximum Link Bandwidth sub-TLV in the BGP-LS Attribute associated with the Link NLRI which carries the corresponding MT-ID. The bandwidth advertised can be exclusive for this logical topology. The advertisement of other topology-specific TE attributes in BGP-LS is for further study. The receiving BGP-LS speaker should be prepared to receive any TE attributes in BGP-LS Attribute with the associated MT-ID carried in the corresponding Link NLRI.

### 4. Scalability Considerations

This document assumes that each NRP is associated with an independent logical topology, and for the inter-domain NRPs, the MT-IDs used in the involved domains are consistent, or some mapping mechanism is used, so that the associated MT-ID can be used to identify the NRP in the control plane. Using MT-ID to identify NRP allows to reuse the multi-topology related control plane mechanisms for the distribution of NRP topology and resource information, while this 1:1 mapping between NRP and MT-ID also has some limitations. For example, even if multiple NRPs share the same topology, each NRP still need to be

identified using a unique MT-ID in the control plane. Therefore independent path computation needs be executed for each NRP. The number of NRPs supported in a network may be dependent on the number of topologies supported, which is related to both the number of topologies supported in the protocol and the control plane overhead which the network could afford. Since no new control protocol extension is required, the mechanism described in this document is considered useful for network scenarios in which the required number of NRPs is small (e.g., less than 10). For network scenarios where the number of required NRPs is large, more scalable solutions would be needed which may require further protocol extensions and enhancements.

## 5. Security Considerations

The security considerations in [RFC9552] [RFC9085] and [RFC9514] apply to this document.

This document introduces no additional security vulnerabilities to BGP-LS. The mechanism proposed in this document is subject to the same vulnerabilities as any other protocol that relies on BGP-LS.

## 6. IANA Considerations

This document does not request any IANA actions.

## 7. Acknowledgments

The authors would like to thank Shunwan Zhuang, Adrian Farrel, Susan Hares, Jeffrey Haas, Ketan Talaulikar and Mohamed Boucadair for the review and discussion of this document.

## 8. References

### 8.1. Normative References

[I-D.ietf-lsr-isis-sr-vtn-mt]

Xie, C., Ma, C., Dong, J., and Z. Li, "Applicability of IS-IS Multi-Topology (MT) for Segment Routing based Network Resource Partition (NRP)", Work in Progress, Internet-Draft, draft-ietf-lsr-isis-sr-vtn-mt-11, 13 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-isis-sr-vtn-mt-11>>.

[I-D.ietf-spring-resource-aware-segments]

Dong, J., Miyasaka, T., Zhu, Y., Qin, F., and Z. Li, "Introducing Resource Awareness to SR Segments", Work in Progress, Internet-Draft, draft-ietf-spring-resource-

aware-segments-15, 3 September 2025,  
<<https://datatracker.ietf.org/doc/html/draft-ietf-spring-resource-aware-segments-15>>.

- [I-D.ietf-spring-sr-for-enhanced-vpn]  
Dong, J., Miyasaka, T., Zhu, Y., Qin, F., and Z. Li,  
"Segment Routing based Network Resource Partition (NRP)  
for Enhanced VPN", Work in Progress, Internet-Draft,  
draft-ietf-spring-sr-for-enhanced-vpn-09, 10 June 2025,  
<<https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-for-enhanced-vpn-09>>.
- [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>>.
- [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>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S.,  
Decraene, B., Litkowski, S., and R. Shakir, "Segment  
Routing with the MPLS Data Plane", RFC 8660,  
DOI 10.17487/RFC8660, December 2019,  
<<https://www.rfc-editor.org/info/rfc8660>>.
- [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>>.
- [RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler,  
H., and M. Chen, "Border Gateway Protocol - Link State  
(BGP-LS) Extensions for Segment Routing", RFC 9085,  
DOI 10.17487/RFC9085, August 2021,  
<<https://www.rfc-editor.org/info/rfc9085>>.
- [RFC9086] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Patel, K.,  
Ray, S., and J. Dong, "Border Gateway Protocol - Link  
State (BGP-LS) Extensions for Segment Routing BGP Egress  
Peer Engineering", RFC 9086, DOI 10.17487/RFC9086, August  
2021, <<https://www.rfc-editor.org/info/rfc9086>>.



- [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>>.
- [RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <<https://www.rfc-editor.org/info/rfc9543>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

## 8.2. Informative References

- [RFC9732] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for NRP-Based Enhanced Virtual Private Networks", RFC 9732, DOI 10.17487/RFC9732, March 2025, <<https://www.rfc-editor.org/info/rfc9732>>.

## Authors' Addresses

Chongfeng Xie  
China Telecom  
China Telecom Beijing Information Science & Technology, Beiqijia  
Beijing  
102209  
China  
Email: [xiechf@chinatelecom.cn](mailto:xiechf@chinatelecom.cn)

Cong Li  
China Telecom  
China Telecom Beijing Information Science & Technology, Beiqijia  
Beijing  
102209  
China  
Email: [licong@chinatelecom.cn](mailto:licong@chinatelecom.cn)

Jie Dong  
Huawei Technologies  
Huawei Campus, No. 156 Beiqing Road  
Beijing  
100095  
China  
Email: jie.dong@huawei.com

Zhenbin Li  
Huawei Technologies  
Huawei Campus, No. 156 Beiqing Road  
Beijing  
100095  
China  
Email: lizhenbin@huawei.com