

IDR Working Group  
Internet-Draft  
Intended status: Informational  
Expires: 15 December 2025

C. Xie  
C. Li  
China Telecom  
J. Dong  
Z. Li  
Huawei Technologies  
13 June 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-11

## Abstract

An NRP is a subset of the network resources and associated policies on each of a connected set of links in the underlay network. An NRP could be used as the underlay to support one or a group of enhanced VPN services, which provide advanced features such as guaranteed resources, bounded latency or jitter to meet specific customer connectivity requirements.

When Segment Routing (SR) is used for building 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. In some network scenarios, each NRP can be associated with a unique logical network topology. When a centralized network controller is used for NRP-specific constraint-based path computation, especially when an NRP spans multiple IGP areas or multiple Autonomous Systems (ASes), BGP-LS is needed to advertise the NRP topology and associated resource information to the network controller. This document describes a mechanism to distribute the information of SR based NRPs using BGP-LS with MT.

## 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 December 2025.

## 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 Attribute for SR-based NRP . . . . .	4
2.1. Intra-domain Topology Advertisement . . . . .	4
2.2. Inter-Domain Topology Advertisement . . . . .	5
3. Advertisement of Resource Attribute for SR-based NRP . . . . .	6
4. Scalability Considerations . . . . .	7
5. Security Considerations . . . . .	7
6. IANA Considerations . . . . .	8
7. Acknowledgments . . . . .	8
8. References . . . . .	8
8.1. Normative References . . . . .	8
8.2. Informative References . . . . .	9
Authors' Addresses . . . . .	9

## 1. Introduction

Enhanced VPNs aim to deliver VPN services with enhanced characteristics to customers who have specific requirements on their connectivity, such as guaranteed resources, and bounded latency or jitter. Enhanced VPNs require integration between the overlay VPN connectivity and the characteristics provided by the underlay network. [RFC9543] discusses the general framework, components, and interfaces for requesting and operating network slices using IETF technologies. Providing network slice services is considered as one target use case of enhanced VPNs.

[RFC9543] also introduces the concept of the Network Resource Partition (NRP), which is 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 the 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 requirement of one or a group of enhanced VPN services. The enforcing mechanism for NRP is outside the scope of this document. To meet the requirement of enhanced VPN services, a number of NRPs can be created, each with a subset of network resources allocated on network nodes and links in a customized topology of the physical network. The exact mechanism to map one or group of enhanced VPN services to a specific NRP is outside the scope of this document.

[I-D.ietf-spring-resource-aware-segments] introduces resource awareness to Segment Routing (SR) [RFC8402]. The resource-aware Segment Identifiers (SIDs) have additional semantics to identify the set of network resources available for the packet processing action associated with the SIDs. As described in [I-D.ietf-spring-sr-for-enhanced-vpn], the resource-aware SIDs can be used to build SR-based NRPs with the required network topology and network resource attributes to support enhanced VPN services. With SR-based data plane, SIDs can be used to represent both the topological instructions and a subset of network resources on the network nodes and links which are allocated to an NRP.

To allow NRP-specific constraint-based path computation and/or NRP-specific shortest path computation to be performed by network controller and network nodes, the set of resource-aware SR SIDs and the associated topology and resource attributes of an NRP need to be distributed using a control plane. When a centralized network controller is used for NRP-specific constraint-based path computation, especially when an NRP spans multiple IGP areas or multiple Autonomous Systems (ASes), BGP-Link State (BGP-LS) [RFC9552] is needed to advertise the NRP information in each IGP area or AS to the network controller, so that the controller could use the collected information to build the view of inter-area or inter-AS SR NRPs.

In some network scenarios, the required number of NRPs could be small, and it can be assumed that each NRP is associated with an independent topology and has a set of dedicated or shared network resources. [I-D.ietf-lsr-isis-sr-vtn-mt] describes the IS-IS Multi-Topology (MT) [RFC5120] based mechanism to advertise an independent topology and the associated SR SIDs, together with the resource and

Traffic Engineering (TE) attributes for each SR based NRP. This document describes a mechanism to distribute the information of SR based NRPs to the network controller using BGP-LS with Multi-Topology.

## 2. Advertisement of Topology Attribute for SR-based NRP

[I-D.ietf-lsr-isis-sr-vtn-mt] describes the IS-IS Multi-Topology based mechanisms to distribute the topology and the SR SIDs associated with SR based NRPs. This section describes the corresponding BGP-LS mechanism to distribute both the intra-domain and inter-domain topology attributes of SR based NRPs. It is considered that in each domain, one data plane mechanism is used for one NRP, while for inter-domain SR based NRPs, different data plane mechanisms may be used in different domains. For the inter-domain SR based NRPs, 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 (Type 263), which can contain one or more IS-IS or OSPF Multi-Topology Identifiers for a link, node, or prefix. The rules of the usage of MT-ID TLV is described in section 5.2.2.1 of [RFC9552] as follows:

"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). When included as a Link or Prefix Descriptor, only a single MT-ID TLV containing the MT-ID of the topology where the link or the prefix is reachable is allowed. In case one wants to advertise multiple topologies for a given Link or Prefix Descriptor, multiple NLRIs MUST be generated where each NLRI contains a single unique MT-ID."

[RFC9085] defines the BGP-LS extensions to carry the SR-MPLS information using TLVs of BGP-LS Attribute. When Multi-Topology is used with the SR-MPLS data plane, 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, the MT-ID TLV carried in the prefix descriptor or link descriptor [RFC9552] can be used to identify the corresponding topology of the SIDs.

[RFC9514] defines the BGP-LS extensions to advertise Segment Routing over IPv6 (SRv6) information along with their functions and attributes. When Multi-Topology is used with the SRv6 data plane,

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. 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 Egress Peer Engineering with SR-MPLS. The BGP-LS extensions for Egress Peering Engineering with SRv6 are specified in [RFC9514]. Such information could be used by a network controller for the computation and instantiation of inter-AS SR-TE paths.

In some network scenarios, for instance, an operator's network consists of multiple parts, such as metro area networks, backbone networks, or data center networks, each part being a different AS. Thus there is a need to create NRPs which span multiple ASes. The inter-domain NRPs may have different inter-domain connectivity, and may be associated with different subsets of network resources in each domain and also on the inter-domain links. To build multi-domain SR based NRPs, the inter-domain topology and the associated BGP Peering SIDs of each NRP for the inter-domain links need to be advertised.

When MT-ID is used consistently in multiple domains covered by an NRP, the topology-specific BGP peering SIDs can be advertised with the MT-ID carried in the corresponding Link NLRI. This can be achieved with the existing mechanisms as defined in [RFC9552][RFC9086] and [RFC9514].

Depending on the requirement of inter-domain NRPs, different mechanisms can be used on the inter-domain connection:

- \* One External BGP (EBGP) session between two ASes can be established over multiple underlying links. In this case, different underlying links can be used for different inter-domain NRPs, which requires the links to be isolated from each other. In another similar case, the EBGP session is established over a single link, while the network resource (e.g. bandwidth) on this link can be partitioned into several pieces, each of which can be considered as a virtual member link. An NRP can be associated with one of the underlying physical or virtual member links. In both cases, different BGP Peer-Adj-SIDs or SRv6 End.X SIDs need to be allocated to each underlying physical or virtual member link, and the association between the BGP Peer-Adj-SID/End.X SID and the MT-ID of the NRP needs to be advertised by the ASBR.
- \* For inter-domain connection between two ASes, multiple EBGP sessions can be established between different sets of peering ASBRs. It is possible that some of these BGP sessions are used for one inter-domain NRP, while some other BGP sessions are used for another inter-domain NRP. In this case, different BGP Peer Node SIDs need to be allocated to each BGP session and are advertised using the mechanism in [RFC9086] and [RFC9514]. The association between the BGP Peer Node SIDs and the MT-ID of the NRP also needs to be advertised by the ASBR.
- \* At the AS-level topology, different inter-domain NRPs may have different inter-AS connectivity. In this case, different BGP Peer Set SIDs need to be allocated to represent the groups of BGP peers which can be used for load-balancing in each inter-domain NRP. The association between the BGP Peer Set SIDs and the MT-ID of the NRP needs to be advertised by the ASBR.

In network scenarios where consistent allocation of MT-ID among multiple domains can not be achieved, the MT-ID advertised by the peering ASBRs of an inter-domain link could be different. Some mapping mechanism may be used by the controller to match the MT-IDs of an inter-domain link in two directions, and concatenate the inter-domain topology of the NRP. Alternatively, a globally-significant NRP identifier may be introduced to identify the inter-domain links of an NRP. Within each domain, the MT based mechanism could be reused for intra-domain topology advertisement. The detailed mechanism is out of the scope of this document.

### 3. Advertisement of Resource Attribute for SR-based NRP

[I-D.ietf-lsr-isis-sr-vtn-mt] specifies the mechanism to advertise the resource and TE attributes associated with each NRP. This section describes the corresponding BGP-LS mechanisms for reporting NRP resource and TE attributes to network controllers.

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

When the Maximum Link Bandwidth sub-TLV is carried in the BGP-LS Attribute associated with the Link NLRI of an NRP, it indicates the amount of link bandwidth resource allocated to the corresponding NRP on the link. The bandwidth allocated to an NRP can be exclusive for traffic in the corresponding NRP. The advertisement of other TE attributes in BGP-LS for NRP is for further study. The receiving BGP-LS speaker MUST be prepared to receive any TE attributes for NRP in BGP-LS Attribute with the associated MT-ID carried in the corresponding Link NLRI.

#### 4. Scalability Considerations

The mechanism described in this document assumes that each NRP is associated with an independent topology, and for the inter-domain NRPs, the MT-IDs used in the involved domains are consistent, so that the MT-IDs can be reused to identify the NRPs in the control plane. Reusing MT-ID can avoid introducing new mechanisms with similar functionality in the control plane, while it 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. Thus 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 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. A detailed analysis about the NRP scalability and the possible optimizations for supporting a large number of NRPs is described in [I-D.ietf-teas-nrp-scalability].

#### 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 and Jeffrey Haas for the review and discussion of this document.

## 8. References

### 8.1. Normative References

- [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-12, 8 May 2025,  
<[https://datatracker.ietf.org/doc/html/draft-ietf-spring-  
resource-aware-segments-12](https://datatracker.ietf.org/doc/html/draft-ietf-spring-resource-aware-segments-12)>.
- [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-08, 12 October 2024,  
<[https://datatracker.ietf.org/doc/html/draft-ietf-spring-  
sr-for-enhanced-vpn-08](https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-for-enhanced-vpn-08)>.
- [RFC8402] Filtsfils, 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>>.
- [RFC9085] Previdi, S., Talaulikar, K., Ed., Filtsfils, 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., Filtsfils, 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>>.
- [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>>.

## 8.2. Informative 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-10, 13 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-isis-sr-vtn-mt-10>>.
- [I-D.ietf-teas-nrp-scalability]  
Dong, J., Li, Z., Gong, L., Yang, G., and G. S. Mishra, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-07, 2 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-07>>.
- [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>>.

## Authors' Addresses

Chongfeng Xie  
China Telecom  
China Telecom Beijing Information Science & Technology, Beiqijia  
Beijing  
102209  
China  
Email: xiechf@chinatelecom.cn

Cong Li  
China Telecom  
China Telecom Beijing Information Science & Technology, Beiqijia  
Beijing  
102209  
China  
Email: 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