Workgroup: IDR Working Group
Internet-Draft:
draft-ietf-idr-bgpls-sr-vtn-mt-02
Published: 13 March 2023
Intended Status: Informational
Expires: 14 September 2023
Authors: C. Xie, China Telecom
C. Li, China Telecom
J. Dong, Huawei Technologies
Z. Li, Huawei Technologies

BGP-LS with Multi-topology for Segment Routing based Virtual Transport Networks

Abstract

Enhanced VPN (VPN+) aims to provide enhanced VPN service to support some applications' needs of enhanced isolation and stringent performance requirements. VPN+ requires integration between the overlay VPN and the underlay network. A Virtual Transport Network (VTN) is a virtual underlay network which consists of a subset of the network topology and network resources allocated from the physical network. A VTN could be used as the underlay for one or a group of VPN+ services.

When Segment Routing is used as the data plane of VTNs, each VTN can be allocated with a group of Segment Identifiers (SIDs) to identify the topology and resource attributes of network segments in the VTN. The association between the network topology, the network resource attributes and the SR SIDs may need to be distributed to a centralized network controller. In network scenarios where each VTN can be associated with a unique logical network topology, this document describes a mechanism to distribute the information of SR based VTNs using BGP-LS with Multi-Topology.

Status of This Memo

Table of Contents

1. Introduction

Enhanced VPN (VPN+) provides enhancement to VPN services to support the needs of new applications, particularly including the applications that are associated with 5G services. These applications require enhanced isolation and stringent performance requirements. VPN+ requires integration between the overlay connectivity and the characteristics provided by the underlay networks. [I-D.ietf-teas-enhanced-vpn] specifies the framework of VPN+ and describes the candidate component technologies in different network planes and layers. VPN+ can be used to underpin network slicing, and will also be of use in more generic scenarios.

To meet the requirement of VPN+ services, a number of Virtual Transport Networks (VTNs) need to be created, each of which consists of a subset of network resources allocated from the underlay network, and is associated with a customized logical topology. A VTN can be used to support one or a group of VPN+ services.

[I-D.ietf-spring-resource-aware-segments] introduces resource awareness to Segment Routing (SR) [RFC8402]. The resource-aware 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 segments can be used to build SR based VTNs with the required network topology and network resource attributes to support VPN+ services.

To allow the VTN-specific constraint-based path computation and/or VTN-specific shortest path computation to be performed by network controller and network nodes, the group of resource-aware SIDs allocated by the network nodes for the VTN, together with the associated topology and resource attributes of the VTN need to be distributed in the control plane. When a centralized network controller is used for VTN-specific constraint-based path computation, especially when a VTN spans multiple IGP areas or multiple Autonomous Systems (ASes), BGP-LS is needed to advertise the VTN 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 VTNs.

In some network scenarios, it is assumed that each VTN 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 IGP Multi-Topology (MT) [RFC5120] based mechanism to advertise the topology and the associated SR SIDs, together with the resource and TE attributes for each SR based VTN. This document describes a mechanism to distribute the information of SR based VTNs to the network controller using BGP-LS [I-D.ietf-idr-rfc7752bis] with Multi-Topology.

2. Advertisement of SR VTN Topology Attribute

[I-D.ietf-lsr-isis-sr-vtn-mt] describes the IS-IS Multi-topology based mechanisms to distribute the topology and the associated SR SIDs of SR based VTNs. This section describes the corresponding BGP-LS mechanism to distribute both the intra-domain and inter-domain topology attributes of SR based VTNs.

2.1. Intra-domain Topology Advertisement

In section 4.2.2.1 of [I-D.ietf-idr-rfc7752bis], Multi-Topology Identifier (MT-ID) TLV is defined, which can contain one or more IS-IS or OSPF Multi-Topology IDs. The MT-ID TLV MAY be present in a Link Descriptor, a Prefix Descriptor, or the BGP-LS Attribute of a Node NLRI.

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

[I-D.ietf-idr-bgpls-srv6-ext] defines the BGP-LS extensions to advertise SRv6 segments along with their functions and attributes. When Multi-Topology is used with 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 in the 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.

[I-D.ietf-idr-rfc7752bis] also defines the rules of the usage of MT-ID TLV:

"In 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 Descriptor or Prefix Descriptor, multiple NLRIs MUST be generated where each NLRI contains a single unique MT-ID."

Editor's note: the above rules indicates that only one MT-ID is allowed to be carried the Link or Prefix descriptors. When a link or prefix needs to be advertised in multiple topologies, multiple NLRIs needs to be generated to report all the topologies the link or prefix participates in, together with the topology-specific segment routing information and link attributes. This may increase the number of BGP Updates needed for advertising MT-specific topology attributes, and may introduce additional processing burden to both the sending BGP speaker and the receiving network controller. When the number of topologies in a network is not a small number, some optimization may be needed for the reporting of multi-topology information and the associated segment routing information in BGP-LS. Based on the WG's opinion, this may be elaborated in a future version.

2.2. Inter-Domain Topology Advertisement

[RFC9086] and [I-D.ietf-idr-bgpls-srv6-ext] defines the BGP-LS extensions for advertisement of BGP inter-domain topology information and the BGP Egress Peering Segment Identifiers. Such information could be used by a network controller for the computation and instantiation of inter-AS SR TE paths.

In some network scenarios, there are needs to create VTNs which span multiple ASes. The inter-domain VTNs could have different inter-domain connectivity, and may be associated with different set of network resources in each domain and also on the inter-domain links. In order to build the multi-domain SR based VTNs, it is necessary to advertise the topology and the associated BGP Peering SIDs of each VTN for inter-domain links.

When MT-ID is used consistently in multiple domains covered by a VTN, 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 [I-D.ietf-idr-rfc7752bis][RFC9086] and [I-D.ietf-idr-bgpls-srv6-ext].

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

* One 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 VTNs which requires link isolation between 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. A VTN 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 SID SHOULD be allocated to each underlying physical or virtual member link, the association between the BGP Peer Adj-SID/End.X SID and the MT-ID of the VTN SHOULD be advertised by the ASBR.

* For inter-domain connection between two ASes, multiple EBGP sessions can be established between different set of peering ASBRs. It is possible that some of these BGP sessions are used for one inter-domain VTN, while some other BGP sessions are used for another inter-domain VTN. In this case, different BGP Peer Node SIDs SHOULD be allocated to each BGP session and are advertised using the mechanism in [RFC9086] and [I-D.ietf-idr-bgpls-srv6-ext], the association between the BGP Peer Node SIDs and the MT-ID of the VTN SHOULD be advertised by the ASBR.

* At the AS-level topology, different inter-domain VTNs may have different inter-AS connectivity. Then different BGP Peer Set SIDs MAY be allocated to represent the groups of BGP peers which can be used for load-balancing in each inter-domain VTN. The association between the BGP Peer Node SIDs and the MT-ID of the VTN SHOULD be advertised by the ASBR.

In network scenarios where consistent usage of MT-ID among multiple domains can not be achieved, a global-significant identifier MAY be introduced to identify the inter-domain topology of a VTN. Within each domain, the MT based mechanism could be reused for intra-domain topology advertisement. The detailed mechanism is specified in [I-D.dong-idr-bgpls-sr-enhanced-vpn].

3. Advertisement of SR VTN Resource Attribute

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

The information of the network resources and TE attributes associated with a link of a VTN can be specified by carrying the TE Link attribute TLVs in BGP-LS Attribute [I-D.ietf-idr-rfc7752bis], 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 a VTN, it indicates the amount of link bandwidth resource allocated to the corresponding VTN on the link. The bandwidth allocated to a VTN can be exclusive for traffic in the corresponding VTN. The advertisement of other TE attributes in BGP-LS for VTN is for further study.

4. Scalability Considerations

The mechanism described in this document requires that each VTN is associated with an independent topology, and for the inter-domain VTNs, the MT-IDs used in all the involved domains need to be consistent. Reusing MT-ID as the identifier of VTN can avoid introducing new mechanism with similar functionality in the control plane, while it also has some limitations. For example, when multiple VTNs have the same topology, each VTN still need to be identified using a unique MT-ID in the control plane, thus independent path computation needs be executed for each VTN, although the result of computation for these VTNs would be the same. The number of VTNs supported in a network may be dependent on the number of topologies supported, which is related to the control plane overhead. The mechanism described in this document is applicable to network scenarios where the number of required VTN is relatively small. A detailed analysis about the VTN scalability and the possible optimizations for supporting a large number of VTNs is described in [I-D.ietf-teas-nrp-scalability].

5. Security Considerations

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 for the review and discussion of this document.

8. References

8.1. Normative References

[I-D.ietf-idr-bgpls-srv6-ext]
Dawra, G., Filsfils, C., Talaulikar, K., Chen, M., Bernier, D., and B. Decraene, "BGP Link State Extensions for SRv6", Work in Progress, Internet-Draft, draft-ietf-idr-bgpls-srv6-ext-14, 17 February 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgpls-srv6-ext-14>.

[I-D.ietf-idr-rfc7752bis]
Talaulikar, K., "Distribution of Link-State and Traffic Engineering Information Using BGP", Work in Progress, Internet-Draft, draft-ietf-idr-rfc7752bis-16, 20 February 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-rfc7752bis-16>.

[I-D.ietf-spring-resource-aware-segments]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Introducing Resource Awareness to SR Segments", Work in Progress, Internet-Draft, draft-ietf-spring-resource-aware-segments-06, 11 October 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-resource-aware-segments-06>.

[I-D.ietf-spring-sr-for-enhanced-vpn]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Segment Routing based Virtual Transport Network (VTN) for Enhanced VPN", Work in Progress, Internet-Draft, draft-ietf-spring-sr-for-enhanced-vpn-04, 11 October 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-for-enhanced-vpn-04>.

[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>.

[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>.

8.2. Informative References

[I-D.dong-idr-bgpls-sr-enhanced-vpn]
Dong, J., Hu, Z., Li, Z., Tang, X., and R. Pang, "BGP-LS Extensions for Scalable Segment Routing based Enhanced VPN", Work in Progress, Internet-Draft, draft-dong-idr-bgpls-sr-enhanced-vpn-04, 4 March 2022, <https://datatracker.ietf.org/doc/html/draft-dong-idr-bgpls-sr-enhanced-vpn-04>.

[I-D.ietf-lsr-isis-sr-vtn-mt] Xie, C., Ma, C., Dong, J., and Z. Li, "Using IS-IS Multi-Topology (MT) for Segment Routing based Virtual Transport Network", Work in Progress, Internet-Draft, draft-ietf-lsr-isis-sr-vtn-mt-04, 13 March 2023, <https://datatracker.ietf.org/api/v1/doc/document/draft-ietf-lsr-isis-sr-vtn-mt/>.

[I-D.ietf-teas-enhanced-vpn] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Network (VPN+)", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-12, 23 January 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-enhanced-vpn-12>.

[I-D.ietf-teas-nrp-scalability]
Dong, J., Li, Z., Gong, L., Yang, G., Guichard, J., Mishra, G. S., Qin, F., Saad, T., and V. P. Beeram, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-01, 24 October 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-01>.

[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