IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2020

C. Xie
C. Li
China Telecom
J. Dong
Z. Li
Huawei Technologies

March 9, 2020

BGP-LS with Multi-topology for Segment Routing based Virtual Transport Networks
draft-xie-idr-bgpls-sr-vtn-mt-00

Abstract

Enhanced VPN (VPN+) as defined in I-D.ietf-teas-enhanced-vpn aims to provide enhanced VPN service to support applications's needs of enhanced isolation and stringent performance requirements. VPN+ requries integration between the overlay VPN and the underlay network. A Virtual Transport Network (VTN) is a virtual network which consists of a subset of the network toplogy and network resources allocated from the underlay network. A VTN could be used as the underlay for one or a group of VPN+ services.

I-D.dong-idr-bgpls-sr-enhanced-vpn defines the BGP-LS extensions to distribute the information of Segment Routing (SR) based VTNs to external entities, such as the network controllers. This document describes a simplified mechanism to distribute the information of SR based VTNs using BGP-LS with Multi-Topology.

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 10, 2020.

Copyright Notice

Copyright (c) 2020 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

Enhanced VPN (VPN+) is an 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 have more stringent performance requirements than that can be provided with traditional overlay VPNs. These properties cannot be met with pure overlay networks, as they require integration between the underlay and the overlay networks. [I-D.ietf-teas-enhanced-vpn] specifies the framework of enhanced VPN

and describes the candidate component technologies in different network planes and layers. An enhanced VPN can be used for 5G transport network slicing, and will also be of use in more generic scenarios.

To meet the requirement of enhanced VPN services, a number of Virtual Transport Networks (VTNs) need to be created, each with a subset of the underlay network topology and a set of network resources allocated to meet the requirement of a specific VPN+ service or a group of VPN+ services.

[I-D.dong-spring-sr-for-enhanced-vpn] specifies how segment routing (SR) [RFC8402] can be used to build virtual transport networks (VTNs) with the required network topology and network resources, which could be used as the underlay of enhanced VPN services. [I-D.dong-lsr-sr-enhanced-vpn] and [I-D.xie-lsr-isis-sr-vtn-mt] specifies the IGP mechanism and extensions to build a set of SR based VTNs. 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 inter-area or inter-AS SR VTNs.

[I-D.dong-idr-bgpls-sr-enhanced-vpn] defines the BGP-LS extensions to distribute the information of Segment Routing (SR) based VTNs to external entities, such as the network controllers, which allows flexible combination of the topology and resource attribute to build customized VTNs. While in some network scenarios, it is assumed that each VTN has an independent topology and a set of dedicated network resources. For such scenarios, this document describes a simplified mechanism to distribute the information of SR based VTNs using BGP-LS with Multi-Topology.

2. Advertisement of SR VTN Topology

[I-D.xie-lsr-isis-sr-vtn-mt] describes the ISIS Multi-topology mechanisms to distribute the topology attributes 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 3.2.1.5 of [RFC7752], 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.

[I-D.ietf-idr-bgp-ls-segment-routing-ext] defines the BGP-LS extensions to carry the segment routing information using TLVs of BGP-LS Attribute. When MTR 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 and 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 MTR 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.

[RFC7752] 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 need to be generated where each NLRI contains an unique MT-ID. In the BGP-LS attribute of a Node NLRI, one MT-ID TLV containing the array of MT-IDs of all topologies where the node is reachable is allowed."

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 participates in multiple topologies, multiple NLRIs needs to be generated to report all the topologies a link or prefix participates in, together with the topology-specific segment routing information. This would increase the number of BGP Updates 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 introduced for the reporting of multi-topology information and the associated segment routing information in BGP-LS. Based on the WG's This will be elaborated in a future version.

2.2. Inter-Domain Topology Advertisement

[I-D.ietf-idr-bgpls-segment-routing-epe] and [I-D.ietf-idr-bgpls-srv6-ext] defines the BGP-LS extensions for advertisement of BGP topology information between ASes and the BGP Peering Segment Identifiers. Such information could be used by a network controller for the computation and instantiation of inter-AS traffic engineering SR 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 VTNs using segment routing, it is necessary to advertise the topology and resource attribute of VTN on the inter-domain links and the associated BGP Peering SIDs.

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

o 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. In both cases, different BGP Peer-Adj-SIDs SHOULD be allocated to each underlying physical or virtual member link, and ASBRs SHOULD advertise the VTN identifier associated with each BGP Peer-Adj-SID.

o 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 multi-domain VTN, while some other BGP sessions are used for another multi-domain VTN. In this case, different BGP peer-node-SIDs are allocated to each BGP session, and ASBRs SHOULD advertise the VTN identifier associated with each BGP Peer-node-SIDs.

o At the AS-level topology, different multi-domain VTNs may have different inter-domain connectivity. Different BGP Peer-Set-SIDs can be allocated to represent the groups of BGP peers which can be used for load-balancing in each multi-domain VTN.

When MT-ID is used consistently in multiple ASes 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 [RFC7752][I-D.ietf-idr-bgpls-segment-routing-epe] and [I-D.ietf-idr-bgpls-srv6-ext].

In network scenarios where consistent usage of MT-ID among multiple ASes can not be expected, then a global-significant VTN ID needs to be introduced to define the AS level topologies. Within each domain, the MT based mechanism could be used for intra-domain topology advertisement. The detailed mechanism is specified in [I-D.dong-idr-bgpls-sr-enhanced-vpn].

3. Advertisement of VTN Resource Attribute

[I-D.xie-lsr-isis-sr-vtn-mt] specifies the mechanism to advertise the resource information associated with each VTN. This section describes the corresponding BGP-LS mechanisms. Two optional approaches are described in the following sections.

3.1. Advertising Topology specific TE attributes

The information of the network resources associated with a VTN can be specified by carrying the Link TE attribute TLVs [RFC7752] in BGP-LS Attribute, with the associated MT-ID carried in the corresponding Link NLRI.

For example, the Maximum Link Bandwidth TLV associated with a MT-ID could be used to specify the link bandwidth allocated to the corresponding VTN on a link.

3.2. Associating VTNs with L2 Bundle Member Links

In some network scenarios, the network resources allocated to different VTNs are instantiated using different physical or virtual members links of a Layer 3 interface. The TE attributes of each member link can be advertised using the mechanism described in [I-D.ietf-idr-bgp-ls-segment-routing-ext] and [I-D.ietf-idr-bgpls-srv6-ext]. In order to further describe the association of each VTN with the corresponding Layer 2 member link, the MT-ID TLV as defined in [RFC7752] SHOULD be carried as a sub-TLV in the L2 Bundle Member Attribute TLV. In each L2 Bundle member Attribute TLV, the MT-ID TLV contains only one MT-ID, which describes the association of a VTN with the corresponding Bundle Member Link Local Identifier.

4. Scalability Considerations

The mechanism described in this document requires that each VTN has an independent topology, and for inter-domain VTNs, the MT-ID used in each involved domain is consistent. While this brings the benefits of simplicity, it also has some limitations. For example, it means that even if multiple VTNs may have the same topology attribute, they would still need to be identified using different MT-IDs in the control plane. This requires that for each VTN, independent path computation would be executed. The number of VTNs supported in a network may be dependent on the number of topologies supported, which is related to the control plane computation overhead.

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.dong-spring-sr-for-enhanced-vpn]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., and Z. Li, "Segment Routing for Resource Partitioned Virtual Networks", draft-dong-spring-sr-for-enhanced-vpn-06 (work in progress), December 2019.

[I-D.ietf-idr-bgp-ls-segment-routing-ext]
Previdi, S., Talaulikar, K., Filsfils, C., Gredler, H., and M. Chen, "BGP Link-State extensions for Segment Routing", draft-ietf-idr-bgp-ls-segment-routing-ext-16 (work in progress), June 2019.

[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-19 (work in progress), May 2019.

[I-D.ietf-idr-bgpls-srv6-ext]
Dawra, G., Filsfils, C., Talaulikar, K., Chen, M., daniel.bernier@bell.ca, d., and B. Decraene, "BGP Link State Extensions for SRv6", draft-ietf-idr-bgpls-srv6-ext-02 (work in progress), January 2020.

[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>.

[RFC5029] Vasseur, JP. and S. Previdi, "Definition of an IS-IS Link Attribute Sub-TLV", RFC 5029, DOI 10.17487/RFC5029, September 2007, <https://www.rfc-editor.org/info/rfc5029>.

[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>.

8.2. Informative References

[I-D.dong-idr-bgpls-sr-enhanced-vpn]
Dong, J. and Z. Hu, "BGP-LS Extensions for Segment Routing based Enhanced VPN", draft-dong-idr-bgpls-sr-enhanced-vpn-00 (work in progress), November 2019.

[I-D.dong-lsr-sr-enhanced-vpn]
Dong, J., Hu, Z., and S. Bryant, "IGP Extensions for Segment Routing based Enhanced VPN", draft-dong-lsr-sr-enhanced-vpn-02 (work in progress), November 2019.

[I-D.ietf-lsr-isis-srv6-extensions]
Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extension to Support Segment Routing over IPv6 Dataplane", draft-ietf-lsr-isis-srv6-extensions-06 (work in progress), March 2020.

[I-D.ietf-teas-enhanced-vpn]
Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Networks (VPN+) Services", draft-ietf-teas-enhanced-vpn-05 (work in progress), February 2020.

[I-D.xie-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", March 2020, <https://tools.ietf.org/html/draft-xie-lsr-isis-sr-vtn-mt>.

[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>.

[RFC8668] Ginsberg, L., Ed., Bashandy, A., Filsfils, C., Nanduri, M., and E. Aries, "Advertising Layer 2 Bundle Member Link Attributes in IS-IS", RFC 8668, DOI 10.17487/RFC8668, December 2019, <https://www.rfc-editor.org/info/rfc8668>.

Authors' Addresses

Chongfeng Xie
China Telecom
China Telecom Beijing Information Science & Technology, Innovation park, Beiqijia
Beijing 102209
China

Email: xiechf@chinatelecom.cn

Cong Li
China Telecom
China Telecom Beijing Information Science & Technology, Innovation park, 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