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

Y. Zhu
China Telecom
J. Dong
Z. Hu
Huawei Technologies

March 9, 2020

BGP-LS with Flex-Algo for Segment Routing based Virtual Transport Networks
draft-zhu-idr-bgpls-sr-vtn-flexalgo-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 Flex-Algo.

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].

Table of Contents

1. Introduction 2
2. Advertisement of VTN Topology Attribute 3
   2.1. Intra-domain Topology Advertisement 3
   2.2. Inter-Domain Topology Advertisement 4
3. Advertisement of VTN Resource Attribute 5
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 8
Authors' Addresses 9

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 other 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.zhu-lsr-isis-sr-vtn-flexalgo] 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. The mechanism described in [I-D.dong-idr-bgpls-sr-enhanced-vpn] 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 those scenarios, this document describes a simplified mechanism to distribute the information of SR based VTNs using BGP-LS with Flex-Algo.

2. Advertisement of VTN Topology Attribute

[I-D.zhu-lsr-isis-sr-vtn-flexalgo] describes the mechanism of using ISIS Flex-Algo to distribute the topology constraints 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

The Flex-Algo definition [I-D.ietf-lsr-flex-algo] can be used to describe the topological constraints for path computation on a network topology. As specified in [I-D.zhu-lsr-isis-sr-vtn-flexalgo], the topology of a VTN can be determined by applying Flex-Algo on a default topology.

BGP-LS extensions for Flex-Algo [I-D.ietf-idr-bgp-ls-flex-algo] provide the mechanisms to advertise the Flex-Algo definition information. BGP-LS extensions for SR-MPLS [I-D.ietf-idr-bgp-ls-segment-routing-ext] and SRv6 [I-D.ietf-idr-bgpls-srv6-ext] provide the mechanism to advertise the algorithm-specific segment routing information.

In[I-D.ietf-idr-bgp-ls-segment-routing-ext], algorithm-specific prefix-SIDs can be advertised in BGP-LS attribute associated with Prefix NLRI.

In [I-D.ietf-idr-bgpls-srv6-ext], algorithm-specific SRv6 Locators can be advertised in the Prefix NLRI with the SRv6 Locator TLV carried in the associated BGP-LS Attribute, and algorithm-specific End.X SID can be advertised in BGP-LS Attribute associated with the corresponding Link NLRI. Other types of SRv6 SIDs can also be algorithm-specific and are advertised using the SRv6 SID NLRI .

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 Flex-Algo is used consistently in multiple ASes covered by a VTN, the topology-specific BGP peering SIDs can be advertised together with the admin-group (color) of the corresponding Flex-Algo in the BGP-LS attribute.

In network scenarios where consistent usage of Flex-Algo among multiple ASes can not be expected, then the global-significant VTN-ID can be used to define the AS level topologies. Within each domain, the Flex-Algo 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.zhu-lsr-isis-sr-vtn-flexalgo] specifies the mechanism to advertise the resource information associated with each VTN. It is based on the extensions to the advertisement of L2 bundle member links information[RFC8668]. This section defines the corresponding BGP-LS extensions.

A new TLVs is defined to specify the attribute flags of either a Layer 3 link or a L2 bundle member link. It can be carried in BGP-LS attribute which is associated with a Link NLRI, or it could be carried as a sub-TLV in the L2 Bundle Member Attribute TLV. The format of the sub-TLV is as below:

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |              Type             |             Length            |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |             Flags             |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Where:

Type: TBD

Length: 4 octets.

Flags: 16-bit flags. This field is consistent with the Flag field in IS-IS Link Attribute sub-TLV in [RFC5029]. In addition to the flags defined in [RFC5029], A new Flag V is defined in this document. When the V flag is set, it indicates this link is a virtual link.

In order to associate different virtual or physical bundle member links with the corresponding VTNs, each member link SHOULD be assigned with a dedicated admin-group or extended admin-group (color), which is included in the definition of the corresponding Flex-Algo. Note that in this case the admin-group or extended admin-group of the Layer 3 interface SHOULD be set to the union of all the admin-groups of the virtual or physical member links. This is to ensure that the Layer 3 link will be included in the constraint-based computation of the corresponding Flex-Algo.

The TE attributes of each Layer 3 link or Layer 2 bundle member link, such as the bandwidth, the adj-SIDs or the End.X SIDs, can be advertised using the mechanism as defined in [I-D.ietf-idr-bgp-ls-segment-routing-ext][I-D.ietf-idr-bgpls-segment-routing-epe] and [I-D.ietf-idr-bgpls-srv6-ext].

4. Scalability Considerations

The mechanism described in this document requires that each VTN maps to an independent Flex-Algo. Even if multiple VTNs share the same topology constraints. While this brings the benefits of simplicity, it also has some limitations. For example, it means that even if multiple VTNs have the same topology constraints, they would still need to be identified using different Flex-Algos in the control plane. Then 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 Flex-Algos supported, which is related to the control plane computation overhead.

Another aspect which may impact the number of VTNs supported is that at most 128 Flex-Algos can be defined in a network.

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

TBD

7. Acknowledgments

The authors would like to thank Shunwan Zhuang and Zhenbin Li 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-flex-algo]
Talaulikar, K., Psenak, P., Zandi, S., and G. Dawra, "Flexible Algorithm Definition Advertisement with BGP Link-State", draft-ietf-idr-bgp-ls-flex-algo-02 (work in progress), January 2020.

[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-flex-algo]
Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", draft-ietf-lsr-flex-algo-06 (work in progress), February 2020.

[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.zhu-lsr-isis-sr-vtn-flexalgo]
Zhu, Y., Dong, J., and Z. Li, "Using Flex-Algo for Segment Routing based VTN", March 2020, <https://tools.ietf.org/html/draft-zhu-lsr-isis-sr-vtn-flexalgo>.

[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

Yongqing Zhu
China Telecom

Email: zhuyq8@chinatelecom.cn

Jie Dong
Huawei Technologies

Email: jie.dong@huawei.com

Zhibo Hu
Huawei Technologies

Email: huzhibo@huawei.com