

Workgroup: LSR Working Group

Internet-Draft:

draft-zhu-lsr-isis-sr-vtn-flexalgo-05

Published: 11 July 2022

Intended Status: Standards Track

Expires: 12 January 2023

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

Using Flex-Algo for Segment Routing (SR) based Virtual Transport Network (VTN)

Abstract

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

The topological constraints of a VTN can be defined using Flex-Algo, a mechanism to provide distributed constraint-path computation. In some network scenarios, each VTN can be associated with a unique Flex-Algo, and the set of network resources allocated to a VTN can be instantiated as layer-2 sub-interfaces or member links of the layer-3 interfaces. This document describes the mechanisms to build the Segment Routing (SR) based VTNs using SR Flex-Algo and IGP L2 bundles with minor extensions. This document updates RFC 8668 by defining a new flag in the Parent L3 Neighbor Descriptor in the L2 Bundle Member Attributes TLV.

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 12 January 2023.

Copyright Notice

Copyright (c) 2022 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](#)
 - 1.1. [Requirements Language](#)
- 2. [Advertisement of SR VTN Topology Attributes](#)
- 3. [Advertisement of SR VTN Resource Attributes](#)
- 4. [Forwarding Plane Operations](#)
- 5. [Scalability Considerations](#)
- 6. [Security Considerations](#)
- 7. [IANA Considerations](#)
- 8. [Acknowledgments](#)
- 9. [References](#)
 - 9.1. [Normative References](#)
 - 9.2. [Informative References](#)
- [Authors' Addresses](#)

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 could be provided with existing overlay VPN techniques. Thus these properties 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 may 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 (VTN) can be created, each with a subset of the underlay network topology and a set of network resources

allocated from the underlay network to meet the requirement of a specific VPN+ service or a group of VPN+ services. Another possible approach is to create a set of point-to-point paths, each with a set of network resource reserved along the path, such paths are called Virtual Transport Paths (VTPs). Although using a set of dedicated VTPs can provide similar characteristics as VTN, it has some scalability issues due to the per-path state in the network.

[[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](#)], the resource-aware segment identifiers (SIDs) can be used to build VTNs with the required network topology and network resource attributes to support VPN+ services. With a segment routing based data plane, SIDs can be used to represent both the topology and the set of network resources allocated by network nodes to a VTN. The SIDs of each VTN together with its associated topology and resource attributes need to be distributed using the control plane.

[[I-D.dong-lsr-sr-enhanced-vpn](#)] defines the IGP mechanisms and extensions to provide scalable SR based VTNs. The mechanism in [[I-D.dong-lsr-sr-enhanced-vpn](#)] allows flexible combination of the topology and resource attribute to provide a relatively large number of VTNs. In some network scenarios, the number of required VTNs may be small, thus a solution to provide a small number of VTNs may also be desired.

This document describes a mechanism to build the SR based VTNs using SR Flex-Algo [[I-D.ietf-lsr-flex-algo](#)] and IGP L2 bundle [[RFC8668](#)] with minor extensions. With this mechanism, each VTN is associated with a unique Flex-Algo, and the set of network resources allocated to the VTN is instantiated using layer-2 sub-interfaces or layer-2 member links of the L3 interfaces. It can provide a relatively small number of VTNs, and can be considered as a transitional solution for the VTN deployment.

This document updates [[RFC8668](#)] by defining a new flag in the Parent L3 Neighbor Descriptor in the L2 Bundle Member Attributes TLV. [[RFC8668](#)] states that all bit fields not defined in that document "MUST be set to zero on transmission and ignored on receipt". Section 3 of this document defines a new flag and specifies both when it is set and how it should be processed.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in

BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2. Advertisement of SR VTN Topology Attributes

[[I-D.ietf-lsr-flex-algo](#)] specifies the mechanism to provide distributed constraint-path computation, and the usage of SR-MPLS prefix-SIDs and SRv6 locators for steering traffic along the constrained paths.

The Flex-Algo Definition (FAD) is the combination of calculation-type, metric-type and the topological constraints used for path computation. According to the network nodes' participation of a Flex-Algo, and the rules of including or excluding Admin Groups (i.e. colors) and Shared Risk Link Groups (SRLGs), the topology of a VTN can be described using the associated Flex-Algo. If each VTN is associated with a unique Flex-Algo, the Flex-Algo identifier could be reused as the identifier of the VTN in the control plane.

With the mechanisms defined in [[RFC8667](#)] [[I-D.ietf-lsr-flex-algo](#)], SR-MPLS prefix-SID advertisement can be associated with a specific topology and a specific algorithm, which can be a Flex-Algo. This allows the nodes to use the prefix-SIDs to steer traffic along distributed computed constraint paths according to the associated Flex-Algo in a particular topology.

[[I-D.ietf-lsr-isis-srv6-extensions](#)] specifies the IS-IS extensions to support SRv6 data plane, in which the SRv6 locators advertisement is associated with a topology and a specific algorithm, which can be a Flex-Algo. This allows the nodes to use the SRv6 locators to steer traffic along distributed computed constraint paths according to the associated Flex-Algo in a particular topology. In addition, topology/algorithm specific SRv6 End SIDs and End.X SIDs can be used to enforce traffic over the Loop-Free Alternatives (LFA) computed backup paths.

3. Advertisement of SR VTN Resource Attributes

Each VTN can be allocated a set of dedicated network resources on different network nodes and links.

In order for a network controller or the ingress nodes to perform constraint based path computation for each VTN, the resource attributes of each VTN need to be advertised. This way, the network controller or the ingress node can compute an SR TE path in a VTN by taking both the Flex-Algo constraints and the resource attributes of the VTN into consideration.

IS-IS L2 Bundle [[RFC8668](#)] was defined to advertise the link attributes of the layer-2 bundle member links. In this section, it

is extended to advertise the set of network resource attributes associated with different VTNs on a layer-3 link.

The layer-3 link must have the capability of partitioning the link resources into different subsets for the different VTNs it participates in. It may or may not be a bundle of layer-2 links to achieve this. One partition of the link resources can be instantiated as a layer-2 sub-interface, which can be seen as a virtual layer-2 member link of the layer-3 link. If the layer-3 link is a layer-2 link bundle, it is possible that the set of link resources allocated to a specific VTN is provided by one or multiple physical layer-2 member links.

A new flag "E" (Exclusive) is defined in the flag field of the Parent L3 Neighbor Descriptor in the L2 Bundle Member Attributes TLV (25).

```

    0 1 2 3 4 5 6 7
+-+--+--+--+--+--+
|P|E|          |
+-+--+--+--+--+--+
```

E flag: When the E flag is set, it indicates each member link under the Parent L3 link is used exclusively for one VTN, and load sharing among the member links is not allowed. When the E flag is clear, it indicates load balancing and sharing among the member links are allowed.

Note that legacy implementations of [\[RFC8668\]](#) will set the E flag to zero (clear) meaning that load balancing among component links is the default behavior. Further, when a legacy implementation receives an E flag that is set, it will ignore the flag and so will assume that load balancing among component links is allowed even when the sender has requested it to not be used. The Flex-Algo associated with the VTN can be defined that only nodes which support the E flag are included in the constraint-based path computation and packet forwarding of the VTN.

For each virtual or physical layer-2 member link, the Admin Groups (AG) or Extended Admin Group (EAG) attribute MUST be advertised using the mechanisms as defined in [\[RFC8668\]](#). This is for the correlation between the Flex-Algo specific forwarding entries and the layer-2 member link. Other TE attributes as defined in [\[RFC5305\]](#) such as the Maximum Link Bandwidth attribute MAY also be advertised for the constraint-based path computation performed by the controller or the ingress nodes. The SR-MPLS Adj-SIDs or SRv6 End.X SIDs associated with each of the virtual or physical layer-2 member links MUST be advertised according to [\[RFC8668\]](#) and [\[I-D.dong-lsr-l2bundle-srv6\]](#).

In order to correlate the virtual or physical layer-2 member links with the Flex-Algo ID which is used to identify the VTN, each VTN is assigned with a unique Admin Group (AG) or Extended Admin Group (EAG), and the virtual or physical layer-2 member links associated with this VTN are configured with the AG or EAG assigned to the VTN. The AG or EAG of the parent layer-3 link is set to the union of all the AGs or EAGs of its virtual or physical layer-2 member links. In the definition of the Flex-Algo corresponding to the VTN, it MUST use the Include-Any Admin Group rule with only the AG or EAG assigned to the VTN as the link constraints, the Include-All Admin Group rule or the Exclude Admin Group rule MUST NOT be used. This is to ensure that the layer-3 link is included in the Flex-Algo constraint based path computation for each VTN it participates in.

4. Forwarding Plane Operations

For the SR-MPLS data plane, a prefix SID is associated with the paths calculated using the Flex-Algo corresponding to a VTN. An outgoing layer-3 interface is determined for each path. In addition, the prefix-SID also steers the traffic to use the virtual or physical layer-2 member link which is associated with the VTN on the outgoing layer-3 interface for packet forwarding. A forwarding entry MUST be installed in the forwarding plane using the MPLS label that corresponds to the Prefix-SID associated with the Flex-algorithm corresponding to the VTN. The Adj-SIDs associated with the virtual or physical member links of a VTN can be used together with the prefix-SIDs of the same VTN to build SR-MPLS TE paths under the topological and resource constraints of the VTN.

For the SRv6 data plane, an SRv6 Locator is a prefix which is associated with the paths calculated using the Flex-Algo corresponding to a VTN. An outgoing Layer-3 interface is determined for each path. In addition, the SRv6 Locator prefix also steers the traffic to use the virtual or physical layer-2 member link which is associated with the VTN on the outgoing layer-3 interface for packet forwarding. A forwarding entry for the SRv6 Locator prefix MUST be installed in the forwarding plane for the Flex-algorithm corresponding to the VTN. The End.XU SIDs associated with the virtual or physical member links of a VTN can be used together with other types of SRv6 SIDs of the same VTN to build SRv6 paths under the topological and resource constraints of the VTN.

5. Scalability Considerations

The mechanism described in this document assumes that each VTN is associated with a unique Flex-Algo, so that the Flex-Algo IDs can be reused to identify the VTNs in the control plane. While this brings the benefit of simplicity, it also has some limitations. For example, it means that even if multiple VTNs share the same

topological constraints, they still need to be identified using different Flex-Algo IDs in the control plane, then independent path computation needs to be executed for each VTN. The number of VTNs supported in a network may be dependent on the number of Flex-Algos supported, which is related to the number of Flex-Algos supported in the protocol (which is 128) and the control plane overhead on network nodes. 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 can be found in [[I-D.dong-teas-nrp-scalability](#)].

6. Security Considerations

This document introduces no additional security vulnerabilities to IS-IS.

The mechanism proposed in this document is subject to the same vulnerabilities as any other protocol that relies on IGPs.

7. IANA Considerations

This document does not request any IANA actions.

8. Acknowledgments

The authors would like to thank Zhenbin Li, Peter Psenak, Adrian Farrel and Gyan Mishra for the review and discussion of this document.

9. References

9.1. Normative References

[**I-D.dong-lsr-l2bundle-srv6**] Dong, J. and Z. Hu, "Advertising SRv6 SIDs for Layer 2 Bundle Member Links in IGP", Work in Progress, Internet-Draft, draft-dong-lsr-l2bundle-srv6-01, 24 October 2021, <<https://www.ietf.org/archive/id/draft-dong-lsr-l2bundle-srv6-01.txt>>.

[**I-D.ietf-lsr-flex-algo**] Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-20, 18 May 2022, <<https://www.ietf.org/archive/id/draft-ietf-lsr-flex-algo-20.txt>>.

[**I-D.ietf-lsr-isis-srv6-extensions**]

Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over IPv6 Dataplane", Work in Progress, Internet-Draft,

draft-ietf-lsr-isis-srv6-extensions-18, 20 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-lsr-isis-srv6-extensions-18.txt>>.

[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-04, 5 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-04.txt>>.

[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-02, 5 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-sr-for-enhanced-vpn-02.txt>>.

[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+) Services", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-10, 6 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-10.txt>>.

[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>>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[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>>.

[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>>.

9.2. Informative References

[I-D.dong-lsr-sr-enhanced-vpn] Dong, J., Hu, Z., Li, Z., Tang, X., Pang, R., JooHeon, L., and S. Bryant, "IGP Extensions for Scalable Segment Routing based Enhanced VPN", Work in Progress, Internet-Draft, draft-dong-lsr-sr-enhanced-vpn-07, 29 January 2022, <<https://www.ietf.org/archive/id/draft-dong-lsr-sr-enhanced-vpn-07.txt>>.

[I-D.dong-teas-nrp-scalability]

Dong, J., Li, Z., Gong, L., Yang, G., Guichard, J. N., Mishra, G., Qin, F., Saad, T., and V. P. Beeram, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-dong-teas-nrp-scalability-02, 16 May 2022, <<https://www.ietf.org/archive/id/draft-dong-teas-nrp-scalability-02.txt>>.

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