

LSR Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: 5 September 2024

Y. Zhu  
China Telecom  
J. Dong  
Z. Hu  
Huawei Technologies  
4 March 2024

Using Flex-Algo for Segment Routing (SR) based Network Resource  
Partition (NRP)  
draft-zhu-lsr-isis-sr-vtn-flexalgo-07

Abstract

Enhanced VPNs aim to deliver VPN services with enhanced characteristics, such as guaranteed resources, latency, jitter, etc., so as to support customers requirements for connectivity services with these enhanced characteristics. Enhanced VPN requires integration between the overlay VPN connectivity and the characteristics provided by the underlay network. A Network Resource Partition (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.

In some network scenarios, each NRP can be associated with a unique Flexible Algorithm (Flex-Algo), which can provide constraint-path computation based on the customized topological constraints. This document specifies a mechanism to build Segment Routing (SR) based NRPs by combining SR Flex-Algo and IGP L2 bundles with minor extensions.

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 5 September 2024.

## Copyright Notice

Copyright (c) 2024 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
1.1. Requirements Language . . . . .	3
2. Advertisement of NRP Topology Attributes . . . . .	4
3. Advertisement of NRP Resource Attributes . . . . .	4
4. Forwarding Plane Operations . . . . .	6
5. Scalability Considerations . . . . .	7
6. Security Considerations . . . . .	7
7. IANA Considerations . . . . .	7
8. Acknowledgments . . . . .	7
9. References . . . . .	7
9.1. Normative References . . . . .	7
9.2. Informative References . . . . .	9
Authors' Addresses . . . . .	9

## 1. Introduction

Enhanced VPNs aim to deliver VPN services with enhanced characteristics, such as guaranteed resources, latency, jitter, etc., so as to support customers requirements for connectivity services with these enhanced characteristics. Enhanced VPN requires integration between the overlay VPN connectivity and the characteristics provided by the underlay network. [I-D.ietf-teas-ietf-network-slices] discusses the general framework, components, and interfaces for requesting and operating network slices using IETF technologies. Network slice is considered as one target use case of enhanced VPNs.

[I-D.ietf-teas-ietf-network-slices] 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. [I-D.ietf-teas-enhanced-vpn] 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. 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.

[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 SIDs can be used to build SR-based NRPs with the required network topology and network resource attributes to support enhanced VPN services. In an SR-based data plane, Segment Identifiers (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. The SR SIDs and the associated topology and resource attributes of an NRP need to be distributed using a control plane.

In some network scenarios, the required number of NRPs could be small, then it can be assumed that each NRP is associated with an independent logical topology, and has a set of dedicated or shared network resources. For such scenarios, this document describes a simple mechanism to build Segment Routing (SR) based NRPs, by combining SR Flex- Algo and IGP L2 bundles with minor extensions. More specifically, each NRP is associated with a unique Flex- Algo, and the subset of network resources allocated to the NRP are instantiated using either virtual sub-interfaces or layer-2 member links of L3 interfaces.

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 when it should be 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 NRP Topology Attributes

[RFC9350] specifies the mechanism to provide distributed constraint-path computation, and the use of SR-MPLS prefix-SIDs and SRv6 locators in data plane 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 an NRP can be described using the associated Flex-Algo. If each NRP is associated with a unique Flex-Algo, the Flex-Algo identifier could be used as the identifier of the NRP in the control plane.

With the mechanisms defined in[RFC8667] [RFC9350], SR-MPLS prefix-SID advertisement can be associated with a constrained topology, which can be defined by 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.

[RFC9352] specifies the IS-IS extensions to support SRv6 data plane, in which the SRv6 locators advertisement is associated with a constrained topology, which can be defined by a Flex-Algo. This allows the nodes to use the SRv6 locators to steer traffic along constraint paths computed using the Flex-Algo in the associated topology. In addition, SRv6 End SIDs and End.X SIDs which are associated with a Flex-Algo can be used to enforce traffic over the Loop-Free Alternatives (LFA) computed backup paths.

## 3. Advertisement of NRP Resource Attributes

Each NRP can be allocated with a set of dedicated or shared network resources on a connected set of links in the underlay network.

In order for a network controller or the ingress nodes to perform constraint based path computation for each NRP, the resource attributes of each NRP need to be advertised in the control plane, and distributed to the network controller. This way, the network controller or the ingress node can compute an SR-TE path in the NRP by taking both the Flex-Algo constraints and the resource attributes of the NRP 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 different subset of link resources and attributes associated with different NRPs on a layer-3 link.

The layer-3 link must have the capability of partitioning the link resources into different subsets and allocate them to different NRPs it participates in. Each partition of the link resources can be instantiated as a virtual sub-interface, which can be seen as a virtual layer-2 member link of the layer-3 link. When the layer-3 link is a layer-2 link bundle, the subset of link resources allocated to a specific NRP may be provided by one or multiple physical layer-2 member links in the link bundle.

Normally the member links of a L2 link bundle are used in load-balancing mode, some extension is needed to indicate that each member link is used exclusively for traffic of a specific NRP. 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 NRP, and load balancing 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 the 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 NRP can be defined that only nodes which support the E flag and the mechanisms defined in this document are included in the constraint-based path computation and packet forwarding of the NRP.

For each virtual or physical layer-2 member link under the layer-3 interface, the Admin Groups (AG) or Extended Admin Group (EAG) attribute MUST be advertised using the mechanisms as defined in [RFC8668]. Other TE attributes as defined in [RFC5305] such as the Maximum Link Bandwidth attribute MAY be advertised for the constraint-based path computation. The SR-MPLS Adj-SIDs or SRv6 End.X SIDs associated with each of the virtual or physical layer-2 member links SHOULD 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 associated Flex-Algo ID of the NRP, each NRP SHOULD be assigned with a unique Admin Group (AG) or Extended Admin Group (EAG), and the virtual or physical layer-2 member links associated with this NRP SHOULD be configured with the AG or EAG of the NRP. The AG or EAG of the parent layer-3 link SHOULD be 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 which is associated with the NRP, the Include-Any Admin Group with only the AG or EAG assigned to the NRP SHOULD be used as the link constraints, the Include-All Admin Group or the Exclude Admin Group rules MUST NOT be used for a Flex-Algo associated with an NRP. This is to ensure that the layer-3 link is included in the Flex-Algo constraint path computation for each NRP 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 the NRP. 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 NRP 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 NRP. The Adj-SIDs associated with the virtual or physical member links of an NRP MAY be used together with the prefix-SIDs of the same NRP to build SR-MPLS TE paths under the topological and resource constraints of the NRP.

For the SRv6 data plane, an SRv6 Locator is a prefix which is associated with the paths calculated using the Flex-Algo corresponding to the NRP. 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 NRP 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 NRP. The End.XU SIDs associated with the virtual or physical member links of an NRP MAY be used together with other types of SRv6 SIDs of the same NRP to build SRv6 paths under the topological and resource constraints of the NRP.

## 5. Scalability Considerations

The mechanism described in this document assumes that each NRP is associated with a unique Flex-Algo, so that the Flex-Algo IDs can be reused to identify the NRPs in the control plane. While this brings the benefit of simplicity, it also has some limitations. For example, it means that even if multiple NRPs share the same constrained topology, they would still need to be identified using different Flex-Algos in the control plane, then independent path computation needs to be executed for each NRP. Thus 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 nodes could accommodate. The mechanism described in this document is considered useful for network scenarios in which the required number of NRPs is small, as minor control protocol extension is required. For network scenarios where the number of required NRPs is large, more scalable solution 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].

## 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://datatracker.ietf.org/doc/html/draft-dong-lsr-l2bundle-srv6-01>>.

[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-08, 23 October 2023,  
<[https://datatracker.ietf.org/doc/html/draft-ietf-spring-  
resource-aware-segments-08](https://datatracker.ietf.org/doc/html/draft-ietf-spring-resource-aware-segments-08)>.

[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-07, 4 March 2024,  
<[https://datatracker.ietf.org/api/v1/doc/document/draft-  
ietf-spring-sr-for-enhanced-vpn/](https://datatracker.ietf.org/api/v1/doc/document/draft-ietf-spring-sr-for-enhanced-vpn/)>.

[I-D.ietf-teas-enhanced-vpn]

Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A  
Framework for NRP-based Enhanced Virtual Private Network",  
Work in Progress, Internet-Draft, draft-ietf-teas-  
enhanced-vpn-17, 25 December 2023,  
<[https://datatracker.ietf.org/doc/html/draft-ietf-teas-  
enhanced-vpn-17](https://datatracker.ietf.org/doc/html/draft-ietf-teas-enhanced-vpn-17)>.

[I-D.ietf-teas-ietf-network-slices]

Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani,  
K., Contreras, L. M., and J. Tantsura, "A Framework for  
Network Slices in Networks Built from IETF Technologies",  
Work in Progress, Internet-Draft, draft-ietf-teas-ietf-  
network-slices-25, 14 September 2023,  
<[https://datatracker.ietf.org/doc/html/draft-ietf-teas-  
ietf-network-slices-25](https://datatracker.ietf.org/doc/html/draft-ietf-teas-ietf-network-slices-25)>.

[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>>.
- [RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.
- [RFC9352] Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352, February 2023, <<https://www.rfc-editor.org/info/rfc9352>>.

## 9.2. Informative References

- [I-D.ietf-teas-nrp-scalability]  
Dong, J., Li, Z., Gong, L., Yang, G., Mishra, G. S., and F. Qin, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-03, 21 October 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-03>>.

## 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