

LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 25 April 2024

J. Dong
Z. Hu
Z. Li
Huawei Technologies
X. Tang
R. Pang
China Unicom
S. Bryant
University of Surrey
23 October 2023

IGP Extensions for Scalable Segment Routing based Virtual Transport
Network (VTN)
draft-dong-lsr-sr-enhanced-vpn-10

Abstract

Enhanced VPN (VPN+) aims to provide enhanced VPN services 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 to support one or a group of VPN+ services. In the context of network slicing, a VTN could be instantiated as a network resource partition (NRP).

This document specifies the IGP mechanisms with necessary extensions to advertise the associated topology and resource attributes for scalable Segment Routing (SR) based NRPs. Each NRP can have a customized topology and a set of network resources allocated from the physical network. Multiple NRPs may shared the same topology, and multiple NRPs may share the same set of network resources on some network segments. This allows flexible combination of the network topology and network resource attributes to build a relatively large number of NRPs with a relatively small number of logical topologies. A group of resource-aware SIDs and SRv6 Locators can be assigned to each NRP. The proposed mechanism is applicable to both Segment Routing with MPLS data plane (SR-MPLS) and Segment Routing with IPv6 data plane (SRv6). This document also describes the mechanism of using dedicated NRP ID in the data plane instead of the per-NRP resource-aware SIDs and SRv6 Locators to further reduce the control plane and data plane overhead of maintaining a large number of NRPs.

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 25 April 2024.

Copyright Notice

Copyright (c) 2023 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	3
1.1. Requirements Language	4
2. Advertisement of NRP Definition	5
3. Advertisement of NRP Topology Attribute	6
3.1. MTR based Topology Advertisement	6
3.2. Flex-Algo based Topology Advertisement	7
4. Advertisement of NRP Resource Attributes	8
4.1. Option 1: L2 Bundle based Approach	10
4.2. Option 2: Per-NRP Link TE Attributes	11
5. Advertisement of NRP specific Data Plane Identifiers	11
5.1. Advertisement of NRP-specific SR-MPLS SIDs	11
5.2. Advertisement of NRP-specific SRv6 Locators and SIDs	15
5.2.1. NRP-specific SRv6 Locators and End SIDs	15
5.2.2. NRP-specific SRv6 End.X SIDs	18
5.3. Advertisement of Dedicated Data Plane NRP ID	18

6. Security Considerations	19
7. IANA Considerations	19
8. Contributors	20
9. Acknowledgments	20
10. References	20
10.1. Normative References	20
10.2. Informative References	22
Authors' Addresses	22

1. Introduction

Enhanced VPN (VPN+) is an enhancement to VPN services to support the needs of new applications, particularly 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 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 network slicing, and will also be of use in more generic scenarios.

To meet the requirement of different enhanced VPN services, a number of virtual underlay networks need to be created, each with a customized network topology and a set of network resources allocated from the physical network to meet the requirement of one or a group of VPN+ services. Such a virtual underlay network is called Virtual Transport Network (VTN) in [I-D.ietf-teas-enhanced-vpn].

[I-D.ietf-teas-ietf-network-slices] introduces the concept Network Resource Partition (NRP) as a set of network resources that are available to carry traffic and meet the SLOs and SLEs. In order to allocate network resources to an NRP, the NRP is associated with a network topology to define the set of links and nodes. Thus NRP can be seen as an instantiation of VTN in the context of network slicing. For clarity, the rest of this document uses NRP in the description of the proposed mechanisms and protocol extensions.

[I-D.ietf-spring-resource-aware-segments] introduces resource-aware segments by associating existing type of SIDs with network resource attributes (e.g. bandwidth, processing or storage resources). These resource-aware SIDs retain their original functionality, with the additional semantics of identifying the set of network resources available for the packet processing action.[I-D.ietf-spring-sr-for-enhanced-vpn] describes the use of resource-aware segments to build SR based NRPs. To allow the network controller and network nodes to perform NRP-specific explicit path computation and/or shortest path computation, the group of resource-

aware SIDs allocated by network nodes to each NRP and the associated topology and resource attributes need to be distributed using the control plane.

[I-D.ietf-teas-nrp-scalability] analyzes the scalability requirements and the control plane and data plane scalability considerations of NRP. In order to support a relatively large number of NRPs in the network, one proposed approach is to separate the topology and resource attributes of the NRP in control plane, so that the advertisement and processing of each type of attribute could be decoupled. Multiple NRPs may share the same topology, and multiple NRPs may share the same set of network resources on some network segments, while the difference in either the topology or resource attributes makes them different NRPs. This allows flexible combination of network topology and network resource attributes to build a large number of NRPs with a relatively small number of logical topologies.

This document specifies the IGP control plane mechanisms with necessary extensions for scalable SR based NRPs. A group of resource-aware SIDs and SRv6 Locators can be assigned to each NRP. The proposed mechanism is applicable to both segment routing with MPLS data plane (SR-MPLS) and segment routing with IPv6 data plane (SRv6). This document also describes the mechanisms of using dedicated NRP ID in the data plane instead of the per-NRP resource-aware SIDs to further reduce the control plane and data plane overhead of maintaining a large number of NRPs.

In general this approach applies to both IS-IS and OSPF, while the specific protocol extensions and encodings are different. In the current version of this document, the required IS-IS extensions are described. The required OSPF extensions will be described in a future version or in a separate document.

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 Definition

According to [I-D.ietf-teas-ietf-network-slices], an NRP is a collection of network resources allocated in the underlay network, and is associated with a network topology. Thus an NRP can be defined as the combination of a set of network attributes, which include the topology attribute, the resource attributes, and other possible attributes.

The IS-IS Network Resource Partition Definition (NRPD) sub-TLV is used to advertise the definition of an NRP. It is a sub-TLV of the IS-IS Router-Capability TLV 242 as defined in [RFC7981].

The format of IS-IS NRPD sub-TLV is as below:



Where:

- * Type: TBD
- * Length: The length of the value field of the sub-TLV. It is variable dependent on the included sub-TLVs.
- * NRP ID: A domain significant 32-bit identifier which is used to identify an NRP.
- * MT-ID: 16-bit field which indicates the multi-topology identifier as defined in [RFC5120]. The first 4-bit are set to zero.
- * Algorithm: 8-bit identifier which indicates the algorithm which applies to this NRP. It can be either a normal algorithm [RFC8402] or a Flexible Algorithm [RFC9350].

- * Priority: 8-bit priority value between 0 and 255 inclusive that specifies the priority of the advertisement. Numerically greater values are preferred when there are multiple advertisements of NRPD with the same NRP ID.
- * Sub-sub-TLVs: optional sub-sub-TLVs to specify the additional attributes of an NRP. Currently no sub-sub-TLV is defined in this document.

The NRPD Sub-TLV MAY be advertised in an LSP of any number. A node MUST NOT advertise more than one NRPD Sub-TLV for a given NRP ID.

3. Advertisement of NRP Topology Attribute

This section describes the mechanisms used to advertise the topology attribute associated with SR based NRPs. The topology of an NRP can be determined by the MT-ID and/or the algorithm ID included in the NRP definition. Specifically, in SR networks, the topology of an NRP could be specified using two optional approaches.

The first approach is to use Multi-Topology Routing (MTR) [RFC4915] [RFC5120] with the segment routing extensions to advertise the topology associated with the SR based NRPs. Different algorithms MAY be used to further specify the computation algorithm or the metric type used for path computation within the topology. Multiple NRPs MAY be associated with the same <topology, algorithm> tuple, in that case the IGP computation with the <topology, algorithm> tuple can be shared by these NRPs.

The second approach is to use Flex-Algo [RFC9350] to describe the topological constraints of SR based NRPs on a shared network topology (e.g. the default topology). Multiple NRPs MAY be associated with the same Flex-Algo, in that case the IGP computation with this Flex-Algo can be shared by these NRPs.

3.1. MTR based Topology Advertisement

Multi-Topology Routing (MTR) has been defined in [RFC4915] and [RFC5120] to create different network topologies in one network. It also has the capability of specifying customized attributes for each topology. The traditional use cases of multi-topology are to maintain separate topologies for unicast and multicast services respectively, or to create different topologies for IPv4 and IPv6 in a network. There are some limitations when MTR is used with native IP forwarding, the considerations about MT based IP forwarding are described in [RFC5120].

MTR can be used with SR-MPLS data plane. [RFC8667] specifies the IS-IS extensions to support SR-MPLS data plane, in which the Prefix-SID sub-TLVs can be carried in IS-IS TLV 235 (MT IP Reachability) and TLV 237 (MT IPv6 IP Reachability), and the Adj-SID sub-TLVs can be carried in IS-IS TLV 222 (MT-ISN) and TLV 223 (MT IS Neighbor Attribute).

MTR can also be used with SRv6 data plane. [RFC9352] specifies the IS-IS extensions to support SRv6 data plane, in which the MT-ID is carried in the SRv6 Locator TLV. The SRv6 End SIDs are carried as sub-TLVs in the SRv6 Locator TLV, and inherit the topology/algorithm from the parent locator. The SRv6 End.X SIDs are carried as sub-TLVs in the IS-IS TLV 222 (MT-ISN) and TLV 223 (MT IS Neighbor Attribute), and inherit the topology/algorithm from the parent locator.

These IGP extensions for SR-MPLS and SRv6 can be used to advertise and build the topology for a group of SR based NRPs.

An algorithm ID MAY be used to further specify the calculation-type and/or the metric-type used for path computation within the topology.

3.2. Flex-Algo based Topology Advertisement

[RFC9350] specifies the mechanisms to provide distributed computation of constraint-based paths, and how the SR-MPLS prefix-SIDs and SRv6 locators can be used to steer packets along the constraint-based paths.

The Flex-Algo Definition (FAD) can be used to describe the topological constraints for path computation on a network topology. According to the network nodes' participation of a Flex-Algo, the rules of including or excluding specific Administrative Groups (colors), and/or the Shared Risk Link Groups (SRLGs), the topology of an NRP can be determined by applying the associated Flex-Algo on a particular topology (e.g. the default topology).

With the mechanisms defined in [RFC8667] [RFC9350], prefix-SID advertisement can be associated with a <topology, algorithm> tuple, in which the algorithm can be a Flex-Algo. This allows network nodes to use the prefix-SID to steer traffic along the distributed computed paths according to the identified Flex-Algo in the topology.

[RFC9352] specifies the IS-IS extensions to support SRv6 data plane, in which the SRv6 locators advertisement can be associated with a specific topology and a specific algorithm, which can be a Flex-Algo. With the mechanism defined in [RFC9350], The SRv6 locator can be used to steer traffic along distributed computed paths according to the identified Flex-Algo in the topology. In addition, topology/algorithm specific SRv6 End SID and End.X SID can be used to enforce traffic over the LFA computed backup path.

Multiple Flex-Algos MAY be defined to describe the topological constraints on a shared network topology (e.g. the default topology).

4. Advertisement of NRP Resource Attributes

This section specifies the mechanisms to advertise the network resource attributes associated with an NRP. The mechanism of advertising the link resources and attributes associated with NRP is described. The mechanism of advertising node resources and attributes associated with an NRP are for further study.

A new NRP ID Sub-TLV is defined to advertise the NRP ID and the optional TE attributes associated with the NRP on a physical or virtual link. It MAY be carried under the following TLVs

TLV-22 (Extended IS reachability) [RFC5305]

TLV-23 (IS Neighbor Attribute) [RFC5311]

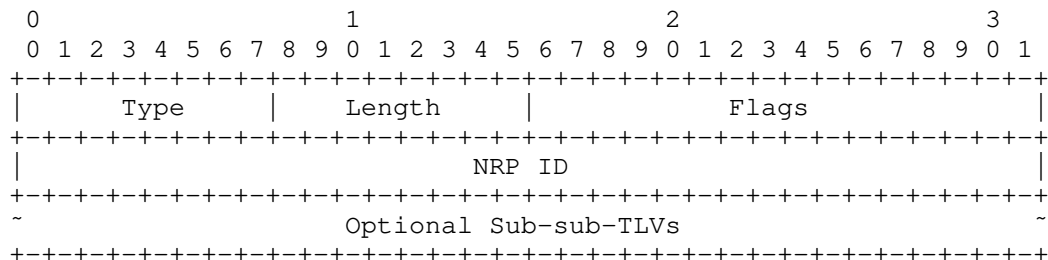
TLV-25 (L2 Bundle Member Attributes TLV) [RFC8668]

TLV-141 (Inter-AS Reachability Information) [RFC5316]

TLV-222 (MT ISN) [RFC5120]

TLV-223 (MT IS Neighbor Attribute) [RFC5311]

The format of the NRP ID Sub-TLV is defined as below:



Where:

- * Type: TBD
- * Length: The length of the value field of the sub-TLV. It is variable dependent on the whether optional Sub-sub-TLVs are included or not.
- * Flags: 16 bit flags. The most significant flag "A" is defined in this document. The rest of the flags SHOULD be set to 0 on transmission and MUST be ignored on receipt.

The format of the Flags field is shown as below:

```

      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
    +---+---+---+---+---+---+---+---+---+
    |A|                                         |
    +---+---+---+---+---+---+---+---+---+

```

A flag: NRP specific Attributes. When the A flag is set to 1, it indicates the NRP ID Sub-TLV further carries NRP-specific TE attributes which are separate from the TE attributes of the associated link. In this case, the optional Sub-sub-TLVs are used to carry the TE attributes associated with this NRP. When the A flag is set to 0, it indicates the NRP inherits the TE attributes of the link.

- * NRP ID: A 32-bit identifier to identify the NRP this member link belongs to.
- * Optional Sub-sub-TLVs: When A flag is set, this field contains the TLVs which are used to advertise the TE attributes associated with this NRP.

Multiple NRP ID Sub-TLVs MAY be advertised for one physical or virtual link to indicate the set of NRPs associated with this link and their optional TE attributes.

For the advertisement of NRP resource attributes, two optional approaches are described in the following sub-sections: the first option is the L2 Bundle [RFC8668] based approach, the second option is to advertise NRP-specific TE attributes directly for the L3 links.

4.1. Option 1: L2 Bundle based Approach

On a Layer 3 interface, each NRP can be allocated with a subset of link resources (e.g. bandwidth). Each subset of link resource can be instantiated as a physical or virtual layer-2 member link under the Layer-3 interface, and the Layer-3 interface is considered as a virtual Layer-2 bundle. The Layer-3 interface may also be a physical Layer 2 link bundle, in this case the subset of link resources allocated to an NRP may be provided by one of the physical Layer-2 member links.

[RFC8668] describes the IS-IS extensions to advertise the link attributes of the Layer 2 member links which comprise a Layer 3 interface. Such mechanism can be extended to advertise the attributes of each physical or virtual member links, and its associated NRPs.

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 are used exclusively for one or a specific group of NRPs, 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.

The NRP ID Sub-TLV defined in this document SHOULD be carried under the L2 Bundle Attribute Descriptors to describe the mapping relationship between the NRPs and the virtual or physical layer 2 member links. As one or more NRPs may share the same set of link resource on a specific link, multiple NRP ID sub-TLVs MAY be advertised under the same virtual or physical member link.

Since Each physical or virtual member link could be associated with a different group of NRPs, each L2 Bundle Attribute Descriptor may carry the link local identifier and attributes of only one Layer 2 member link. Thus multiple L2 Bundle Attribute Descriptors will be used to carry the attributes and the associated NRP IDs of all the Layer 2 member links.

The TE attributes of each virtual or physical member link, such as the bandwidth attributes and the SR SIDs, can be advertised using the mechanism as defined in [RFC8668].

4.2. Option 2: Per-NRP Link TE Attributes

In another case, a Layer-3 interface can participate in multiple NRPs, each of which is allocated with a subset of the forwarding resources of the interface, which can be represented as separate logical data-channels. For each NRP, the TE attributes of the associated data channel can be advertised using NRP-ID Sub-TLV with NRP-specific TE attributes.

The NRP ID Sub-TLV can be used to advertise the link TE attributes associated with each NRP. The existing link TE attributes Sub-TLVs (e.g. Maximum link bandwidth, etc.) can be carried as sub-sub-TLVs under the NRP ID Sub-TLV. In this case the A flag in the NRP ID Sub-TLV SHOULD be set to 1.

5. Advertisement of NRP specific Data Plane Identifiers

In order to steer packets to the NRP-specific paths which are computed taking the topology and network resources of the NRP as the constraints, some fields in the data packet needs to be used to infer or identify the NRP the packet belongs to. As multiple NRPs may share the same topology or Flex-Algo, the topology/Flex-Algo specific SR SIDs or Locators cannot be used to distinguish the packets which belong to different NRPs. Some additional data plane identifiers would be needed to identify the NRP a packet belongs to.

This section describes the mechanisms to advertise the NRP identifiers in different data plane encapsulations.

5.1. Advertisement of NRP-specific SR-MPLS SIDs

With SR-MPLS data plane, the NRP identification information can be implicitly carried in the NRP-specific SIDs. Each node SHOULD allocate a unique Prefix-SID for each NRP it participates in. On a Layer-3 interface, if each Layer 2 member link is associated with only one NRP, the adj-SIDs of the L2 member links could also identify the NRPs. If a member link is associated with multiple NRPs, NRP-specific adj-SIDs MAY need to be allocated to help the NRP-specific local protection.

A new NRP-specific prefix-SID sub-TLV is defined to advertise the prefix-SID and its associated NRP. This sub-TLV MAY be advertised as a sub-TLV of the following TLVs:


TLV-135 (Extended IPv4 Reachability) defined in [RFC5305].

TLV-235 (MT IP Reachability) defined in [RFC5120].

TLV-236 (IPv6 IP Reachability) defined in [RFC5308].

TLV-237 (MT IPv6 IP Reachability) defined in [RFC5120].

The format of the sub-TLV is shown as below:



Where:

- * Type: TBD
- * Length: The length of the value field of the sub-TLV. It is variable dependent on the length of the SID/Index/Label field.
- * Flags: 16-bit flags. The high-order 8 bits are the same as in the Prefix-SID sub-TLV defined in [RFC8667]. The lower-order 8 bits are reserved for future use, which SHOULD be set to 0 on transmission and MUST be ignored on receipt.
- * NRP ID: A 32-bit identifier to identify the NRP this prefix-SID associates with.
- * SID/Index/Label: The same as defined in [RFC8667].

One or more of NRP-specific Prefix-SID sub-TLVs MAY be carried in the Multi-topology IP Reachability TLVs (TLV 235 or TLV 237), the MT-ID of the TLV SHOULD be the same as the MT-ID in the definition of these NRPs.

A new NRP-specific Adj-SID sub-TLV is defined to advertise the adj-SID and its associated NRP. This sub-TLV may be advertised as a sub-TLV of the following TLVs:

TLV-22 (Extended IS reachability) [RFC5305]

TLV-23 (IS Neighbor Attribute) [RFC5311]


TLV-25 (L2 Bundle Member Attributes) [RFC8668]

TLV-141 (Inter-AS Reachability Information) [RFC5316]

TLV-222 (MT ISN) [RFC5120]

TLV-223 (MT IS Neighbor Attribute) [RFC5311]

The format of the sub-TLV is shown as below:



Where:

- * Type: TBD
- * Length: The length of the value field of the sub-TLV. It is variable dependent on the length of the SID/Index/Label field.
- * Flags: 16-bit flags. The high-order 8 bits are the same as in the Adj-SID sub-TLV defined in [RFC8667]. The lower-order 8 bits are reserved for future use, which SHOULD be set to 0 on transmission and MUST be ignored on receipt.
- * NRP ID: A 32-bit global identifier to identify the NRP this Adj-SID associates with.
- * SID/Index/Label: The same as defined in [RFC8667].

One or more NRP-specific Adj-SID sub-TLV MAY be carried in the Multi-topology ISN or Multi-topology IS Attribute TLVs (TLV 222 or TLV 223), the MT-ID of the TLV SHOULD be the same as the MT-ID in the definition of these NRPs.

A new NRP-specific LAN Adj-SID sub-TLV is defined to advertise the adj-SID and its associated NRP for each neighbor on a LAN interface. This sub-TLV may be advertised as a sub-TLV of the following TLVs:

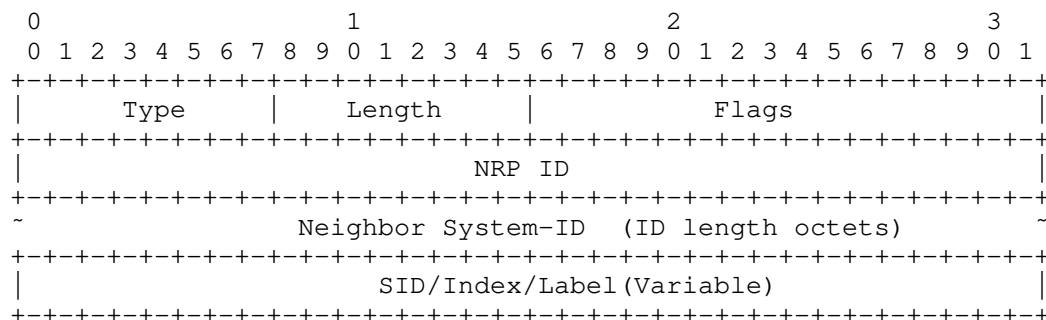
TLV-22 (Extended IS reachability) [RFC5305]

TLV-23 (IS Neighbor Attribute) [RFC5311]

TLV-222 (MT ISN) [RFC5120]

TLV-223 (MT IS Neighbor Attribute) [RFC5311]

The format of the sub-TLV is shown as below:



Where:

- * Type: TBD
- * Length: The length of the value field of the sub-TLV. It is variable dependent on the length of the SID/Index/Label field.
- * Flags: 16-bit flags. The high-order 8 bits are the same as in the Adj-SID sub-TLV defined in [RFC8667]. The lower-order 8 bits are reserved for future use, which SHOULD be set to 0 on transmission and MUST be ignored on receipt.
- * NRP ID: A 32-bit global identifier to identify the NRP this Adj-SID associates with.
- * Neighbor System-ID: IS-IS System-ID of length "ID Length" as defined in [ISO10589].
- * SID/Index/Label: The same as defined in [RFC8667].


One or more NRP-specific LAN Adj-SID sub-TLV MAY be carried in the Multi-topology ISN or Multi-topology IS Attribute TLVs (TLV 222 or TLV 223), the MT-ID of the TLV SHOULD be the same as the MT-ID in the definition of these NRPs.

5.2. Advertisement of NRP-specific SRv6 Locators and SIDs

5.2.1. NRP-specific SRv6 Locators and End SIDs

With SRv6 data plane, the NRP identification information can be implicitly or explicitly carried in the SRv6 Locator of the corresponding NRP, this is to ensure that all network nodes (including both the end nodes and the transit nodes) can identify the NRP to which a packet belongs to. Network nodes SHOULD allocate NRP-specific Locators for each NRP it participates in. The NRP-specific Locators are used as the covering prefix of NRP-specific SRv6 End SIDs, End.X SIDs and other types of SIDs.

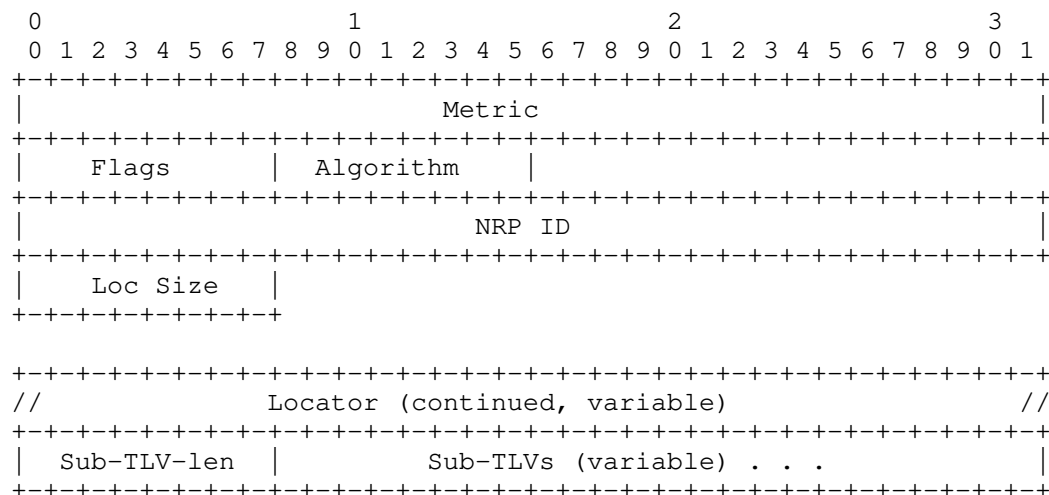
In one possible approach, each NRP-specific Locator is advertised in a separate TLV called "NRP specific SRv6 Locator TLV". Its format is shown as below:



Where:

- * Type: TBD
- * The semantics of the Length field, the R bits and the MT ID field are the same as those defined in [RFC9352].

Followed by one or more locator entries of the form:



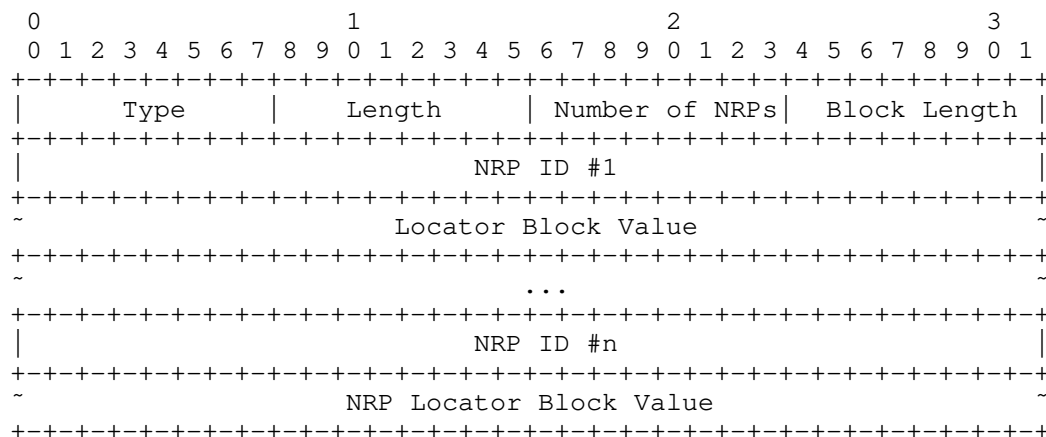
Where:

- * NRP ID: A 32-bit identifier to identify the NRP this Locator is associated with.
- * All the other fields are the same as those defined in [RFC9352].

The NRP-specific SRv6 End SIDs are carried in the NRP-specific SRv6 Locator TLV, and inherits the topology, algorithm and NRP from the parent NRP-specific Locator.

In another possible approach, when a group of NRPs share the same topology/algorithm, the topology/algorithm specific Locator is the covering prefix of a group of NRP-specific Locators. Then the advertisement of NRP-specific locators can be optimized to reduce the amount of Locator TLVs advertised in the control plane.

A new NRP locator-block sub-TLV under the SRv6 Locator TLV is defined to advertise a set of sub-blocks which follows the topology/algorithm specific Locator. Each NRP locator-block value is assigned to one of the NRPs which share the same topology/algorithm.



Where:


- * Type: TBD
- * Length: The length of the value field of the sub-TLV. It is variable dependent on the number of NRPs and the Block Length.
- * Number of NRPs: The number of NRPs which share the same topology/algorithm specific Locator as the covering prefix.
- * Block Length: The length of the NRP locator-block which follows the length of the topology/algorithm specific Locator.
- * NRP ID: A 32-bit identifier to identify the NRP the locator-block is associates with.
- * NRP Locator Block Value: The value of the NRP locator-block for each NRP.

With the NRP locator-block sub-TLV, the NRP-specific Locator can be obtained by concatenating the topology/algorithm specific locator and the NRP locator-block value advertised for the NRP.

The NRP-specific SRv6 End SIDs inherit the topology, algorithm and the NRP from the parent NRP-specific Locator.

5.2.2. NRP-specific SRv6 End.X SIDs

The SRv6 End.X SIDs are advertised as sub-TLVs of TLV 22, 23, 25, 141, 222, and 223. In order to distinguish the End.X SIDs which belong to different NRPs, a new "NRP ID sub-sub-TLV" is introduced under the SRv6 End.X SID sub-TLV and SRv6 LAN End.X SID sub-TLV defined in [RFC9352]. Its format is shown as below:



Where:

- * Type: TBD.
- * Length: the length of the Value field of the TLV. It is set to 4.
- * NRP ID: A 32-bit identifier to identify the NRP this End.X SID associates with.

5.3. Advertisement of Dedicated Data Plane NRP ID

As the number of NRPs increases, with the mechanism described in [I-D.ietf-spring-sr-for-enhanced-vpn], the number of SR SIDs and SRv6 Locators allocated for different NRPs would also increase. In network scenarios where the number of SIDs or Locators becomes a concern, some data plane optimization may be needed to reduce the amount of SR SIDs and Locators allocated. As described in [I-D.ietf-teas-nrp-scalability], one approach is to decouple the data plane identifiers used for topology based forwarding and the identifiers used for the NRP-specific processing. Thus a dedicated data plane NRP ID could be encapsulated in the packet. One possible encapsulation of NRP ID in IPv6 data plane is proposed in [I-D.ietf-6man-enhanced-vpn-vtn-id]. One possible encapsulation of NRP ID in MPLS data plane is proposed in [I-D.li-mpls-enhanced-vpn-vtn-id].

In that case, the NRP ID encapsulated in data plane can have the same value as the NRP ID in control plane, so that the overhead of advertising the mapping between the control plane NRP IDs and the corresponding data plane identifiers could be saved.

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

IANA is requested to assign a new code point in the "sub-TLVs for TLV 242 registry".

Type: TBD1

Description: Network Resource Partition Definition

IANA is requested to assign four new code points in the "sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223 registry".

Type: TBD2

Description: Network Resource Partition Identifier sub-TLV

Type: TBD3

Description: NRP-specific Adj-SID

Type: TBD4

Description: NRP-specific LAN Adj-SID

IANA is requested to assign two new code points in the "Sub-TLVs for TLVs 27, 135, 235, 236 and 237 registry".

Type: TBD5

Description: NRP-specific Prefix-SID

Type: TBD6

Description: NRP locator-block

IANA is requested to assign a new code point in the "IS-IS TLV Codepoints registry".

Type: TBD7

Description: NRP-specific SRv6 Locator TLV

IANA is requested to assign a new code point in the "sub-sub-TLVs for SRv6 End SID and SRv6 End.X SID registry".

Type: TBD8

Description: NRP ID Sub-sub-TLV

8. Contributors

TBD

9. Acknowledgments

The authors would like to thank Mach Chen, Dean Cheng, Lee JooHeon, Hongjie Yang and Guoqi Xu for their review and discussion of this document.

10. References

10.1. Normative References

[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-07, 31 May 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-resource-aware-segments-07>>.

[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-05, 31 May 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-for-enhanced-vpn-05>>.

[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-14, 28 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-enhanced-vpn-14>>.

[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-02, 2 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-02>>.

- [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>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <<https://www.rfc-editor.org/info/rfc4915>>.
- [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>>.
- [RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", RFC 7981, DOI 10.17487/RFC7981, October 2016, <<https://www.rfc-editor.org/info/rfc7981>>.
- [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>>.

10.2. Informative References

[I-D.ietf-6man-enhanced-vpn-vtn-id]
Dong, J., Li, Z., Xie, C., Ma, C., and G. S. Mishra,
"Carrying Virtual Transport Network (VTN) Information in
IPv6 Extension Header", Work in Progress, Internet-Draft,
draft-ietf-6man-enhanced-vpn-vtn-id-05, 6 July 2023,
<<https://datatracker.ietf.org/doc/html/draft-ietf-6man-enhanced-vpn-vtn-id-05>>.

[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>>.

[I-D.li-mpls-enhanced-vpn-vtn-id]
Li, Z. and J. Dong, "Carrying Virtual Transport Network
(VTN) Information in MPLS Packet", Work in Progress,
Internet-Draft, draft-li-mpls-enhanced-vpn-vtn-id-03, 16
October 2022, <<https://datatracker.ietf.org/doc/html/draft-li-mpls-enhanced-vpn-vtn-id-03>>.

Authors' Addresses

Jie Dong
Huawei Technologies
Email: jie.dong@huawei.com

Zhibo Hu
Huawei Technologies
Email: huzhibo@huawei.com

Zhenbin Li
Huawei Technologies
Email: lizhenbin@huawei.com

Xiongyan Tang
China Unicom
Email: tangxy@chinaunicom.cn

Ran Pang
China Unicom
Email: pangran@chinaunicom.cn

Stewart Bryant
University of Surrey
Email: stewart.bryant@gmail.com