

LSR
Internet-Draft
Intended status: Standards Track
Expires: July 25, 2020

Shaofu. Peng
Ran. Chen
Gregory. Mirsky
ZTE Corporation
January 22, 2020

ISIS Extension to Support Network Slicing over IPv6 Dataplane
draft-peng-lsr-isis-network-slicing-srv6-00

Abstract

[I-D.peng-teas-network-slicing] defines an unified TN-slice identifier, i.e., AII(administrative instance identifier) and related processing combined with Segment Routing technology, for the purpose of unified identification of topology, computing, storage resources, unified partition of L2 and L3 link resources, unified basis of underlay path selection within or between domains, and unified flow steering rule with SR policy color scheme. This document describes the ISIS extensions required to support Packet Network Slicing over IPv6 dataplane.

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 July 25, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Advertising SRv6 Capabilities and Maximum SID Depths	3
3.	Advertising Participated TN-slice	3
4.	TN-slice specific SRv6 SIDs and Reachability	3
5.	Advertising Anycast Property	4
6.	Advertising Locators and End SIDs for specific TN-slice	4
6.1.	SRv6 Locator per TN-slice TLV Format	4
6.2.	SRv6 End SID for specific TN-slice	6
7.	Advertising SRv6 Adjacency SIDs for specific TN-slice	6
7.1.	SRv6 End.X SID per TN-slice sub-TLV	7
7.2.	SRv6 LAN End.X SID per TN-slice sub-TLV	8
8.	SRv6 SID Structure Indication	10
9.	Advertising Endpoint Behaviors	10
10.	Security Considerations	10
11.	Acknowledgements	10
12.	IANA Considerations	10
12.1.	SRv6 Locator per TN-slice TLV	10
12.2.	SRv6 End.X SID per TN-slice and SRv6 LAN End.X SID per TN-slice sub-TLVs	11
13.	Normative References	11
	Authors' Addresses	12

[1.](#) Introduction

For a packet network, network slicing requires the underlying network to support partitioning of the network resources to provide the client with dedicated (private) networking, computing, and storage resources drawn from a shared pool. [[I-D.peng-teas-network-slicing](#)] defines an unified TN-slice identifier, i.e., AII(administrative instance identifier) and related processing combined with Segment Routing technology, for the purpose of unified identification of topology, computing, storage resources, unified partition of L2 and L3 link resources, unified basis of underlay path selection within or between domains, and unified flow steering rule with SR policy color scheme.

[[I-D.zch-lsr-isis-network-slicing](#)] describes the ISIS extensions required to support Packet Network Slicing over SR-MPLS dataplane. This document continues to describe the ISIS extensions required to

support Packet Network Slicing over IPv6 dataplane, it will supplement on the basis of [[I-D.ietf-lsr-isis-srv6-extensions](#)].

For SRv6 case, IPv6 address resource is directly used to represent SID, so that different IPv6 block could be allocated to different TN-slice. There are two possible ways to advertise TN-slice specific IPv6 block:

- o Traditional prefix reachability, for default AII (0) specific IPv6 block.
- o New SRv6 Locator per TN-slice advertisement, for nonzero TN-slice specific IPv6 block.

2. Advertising SRv6 Capabilities and Maximum SID Depths

SRv6 capable router can advertise the SRv6 Capabilities sub-TLV and Maximum SRv6 SID Depths (MSD) as defined in [[I-D.ietf-lsr-isis-srv6-extensions](#)].

3. Advertising Participated TN-slice

SRv6 capable router indicates participated TN-slice by advertising the TN-slice identifier Participation sub-TLV as defined in [[I-D.zch-lsr-isis-network-slicing](#)].

4. TN-slice specific SRv6 SIDs and Reachability

A node is provisioned with TN-slice specific locators for each of the TN-slice identified by AII participated by that node. Each locator is a covering prefix for all SIDs provisioned on that node which have the matching AII.

Locators MUST be advertised in the SRv6 Locator per TN-slice TLV (see [Section 6.1](#)). Forwarding entries for the locators advertised in the SRv6 Locator per TN-slice TLV MUST be installed in the forwarding plane of receiving SRv6 capable routers when the associated AII is supported by the receiving node.

Locators are routable and MAY also be advertised in Prefix Reachability TLVs (236 or 237). But locators associated with non default AII SHOULD NOT be advertised in Prefix Reachability TLVs (236 or 237).

Locators associated with default AII (0) SHOULD be advertised in a Prefix Reachability TLV (236 or 237) so that legacy routers (i.e., routers which do NOT support SRv6) will install a forwarding entry for default AII (0) SRv6 traffic.

In cases where a locator advertisement is received in both in a Prefix Reachability TLV and an SRv6 Locator per TN-slice TLV, the Prefix Reachability advertisement MUST be preferred when installing entries in the forwarding plane. This is to prevent inconsistent forwarding entries on SRv6 capable/SRv6 incapable routers.

TN-slice specific SRv6 SIDs are advertised as sub-TLVs in the SRv6 Locator per TN-slice TLV except for TN-slice specific SRv6 End.X SIDs/LAN End.X SIDs which are associated with a specific Neighbor/Link and are therefore advertised as sub-TLVs in TLVs 22, 23, 222, 223, and 141.

TN-slice specific SRv6 SIDs are not directly routable and MUST NOT be installed in the forwarding plane. Reachability to TN-slice specific SRv6 SIDs depends upon the existence of a covering TN-slice specific locator.

Adherence to the rules defined in this section will assure that TN-slice specific SRv6 SIDs associated with a supported AII will be forwarded correctly, while SRv6 SIDs associated with an unsupported AII will be dropped. NOTE: The drop behavior depends on the absence of a default/summary route covering a given locator.

5. Advertising Anycast Property

The same prefix/SRv6 Locator can be advertised by multiple routers. See A-flag defined in [[I-D.ietf-lsr-isis-srv6-extensions](#)] to advertise the anycast property.

6. Advertising Locators and End SIDs for specific TN-slice

The SRv6 Locator per TN-slice TLV is introduced to advertise SRv6 Locators and End SIDs associated with each locator for specific TN-slice.

This new TLV shares the sub-TLV space defined for TLVs 135, 235, 236 and 237.

6.1. SRv6 Locator per TN-slice TLV Format

The SRv6 Locator per TN-slice TLV has the following format:


```

      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      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

SRv6 Locator per TN-slice format

Followed by one or more locator entries of the form:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Metric                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Flags           |                                     AII ...           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  ... AII         |      Loc Size   | Locator (variable)...           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Sub-tlv-len     |      Sub-TLVs (variable) . . .           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

SRv6 Locator entry per TN-slice format

Type: TBA1

Length: variable.

Locator entry:

Metric: 4 octets. As described in [\[RFC5305\]](#).

Flags: 1 octet. The following flags are defined

```

      0
      0 1 2 3 4 5 6 7
+---+---+---+---+---+---+---+
|D|   Reserved   |
+---+---+---+---+---+---+---+

```

where:

D bit: When the Locator is leaked from level-2 to level-1, the D bit MUST be set. Otherwise, this bit MUST be clear. Locators with the D

bit set MUST NOT be leaked from level-1 to level-2. This is to prevent looping.

The remaining bits are reserved for future use. They SHOULD be set to zero on transmission and MUST be ignored on receipt.

AII: 4 octet. Administrative Instance Identifier, As defined in [[I-D.peng-teas-network-slicing](#)], represented as the TN-slice Identifier.

Loc-Size: 1 octet. Number of bits in the Locator field. (1 - 128)

Locator: 1-16 octets. This field encodes the advertised SRv6 Locator. The Locator is encoded in the minimal number of octets for the given number of bits.

Sub-TLV-length: 1 octet. Number of octets used by sub-TLVs

Optional sub-TLVs.

6.2. SRv6 End SID for specific TN-slice

The SRv6 End SID sub-TLV defined in [[I-D.ietf-lsr-isis-srv6-extensions](#)] can be reused in this document to advertise TN-slice specific SRv6 Segment Identifiers (SID) with Endpoint behaviors.

The SRv6 End SID sub-TLV is advertised in the SRv6 Locator per TN-slice TLV defined in the previous section. SRv6 End SIDs inherit the AII from the parent locator.

The SRv6 End SID MUST be a subnet of the associated Locator. SRv6 End SIDs which are NOT a subnet of the associated locator MUST be ignored.

Multiple SRv6 End SIDs MAY be associated with the same locator. In cases where the number of SRv6 End SID sub-TLVs exceeds the capacity of a single TLV, multiple Locator per TN-slice TLVs for the same TN-slice specific locator MAY be advertised. For a given AII/Locator the AII value MUST be the same in all TLVs. If this restriction is not met all TLVs for that AII/Locator MUST be ignored.

7. Advertising SRv6 Adjacency SIDs for specific TN-slice

This document defines two new sub-TLVs of TLV 22, 23, 222, 223, and 141 - namely "SRv6 End.X SID per TN-slice" and "SRv6 LAN End.X SID per TN-slice".

where:

B-Flag: Backup flag. If set, the End.X SID is eligible for protection (e.g., using IPFRR) as described in [[RFC8355](#)].

S-Flag. Set flag. When set, the S-Flag indicates that the End.X SID refers to a set of adjacencies (and therefore MAY be assigned to other adjacencies as well).

P-Flag. Persistent flag. When set, the P-Flag indicates that the End.X SID is persistently allocated, i.e., the End.X SID value remains consistent across router restart and/or interface flap.

Other bits: MUST be zero when originated and ignored when received.

Weight: 1 octet. The value represents the weight of the End.X SID for the purpose of load balancing. The use of the weight is defined in [[RFC8402](#)].

AII: 4 octet. Administrative Instance Identifier, As defined in [[I-D.peng-teas-network-slicing](#)], represented as TN-slice Identifier.

Endpoint Behavior: 2 octets. As defined in [[I-D.ietf-spring-srv6-network-programming](#)]. Legal behavior values for this sub-TLV are defined in [Section 8](#).

SID: 16 octets. This field encodes the advertised SRv6 SID.

Sub-sub-TLV-length: 1 octet. Number of octets used by sub-sub-TLVs

Note that multiple TLVs for the same neighbor may be required in order to advertise all of the SRv6 End.X SIDs associated with that neighbor.

[7.2.](#) SRv6 LAN End.X SID per TN-slice sub-TLV

This sub-TLV is used to advertise a TN-slice specific SRv6 SID associated with a LAN adjacency. Since the parent TLV is advertising an adjacency to the Designated Intermediate System(DIS) for the LAN, it is necessary to include the System ID of the physical neighbor on the LAN with which the SRv6 SID is associated. Given that a large number of neighbors may exist on a given LAN a large number of SRv6 LAN END.X SID per TN-slice sub-TLVs may be associated with the same LAN. Note that multiple TLVs for the same DIS neighbor may be required in order to advertise all of the TN-slice specific SRv6 End.X SIDs associated with that neighbor.

The SRv6 LAN End.X SID per TN-slice sub-TLV has the following format:

SRv6 LAN End.X SID per TN-slice format

AII: 4 octet. Administrative Instance Identifier, As defined in [I-D.peng-teas-network-slicing], represented as TN-slice Identifier.

Endpoint Behavior: 2 octets. As defined in [\[I-D.ietf-spring-srv6-network-programming\]](#). Legal behavior values for this sub-TLV are defined in [Section 8](#).

SID: 16 octets. This field encodes the advertised SRv6 SID.

Sub-sub-TLV-length: 1 octet. Number of octets used by sub-sub-TLVs.

8. SRv6 SID Structure Indication

The SRv6 SID Structure Sub-Sub-TLV defined in [\[I-D.ietf-lsr-isis-srv6-extensions\]](#) can be reused in this document to indicate the SRv6 SID Structure information. The SRv6 SID Structure Sub-Sub-TLV could be an optional Sub-Sub-TLV of:

SRv6 End SID Sub-TLV ([Section 6.2](#))

SRv6 End.X SID per TN-slice Sub-TLV ([Section 7.1](#))

SRv6 LAN End.X SID per TN-slice Sub-TLV ([Section 7.2](#))

9. Advertising Endpoint Behaviors

The Endpoint behaviors and their codepoints, which MAY be advertised by IS-IS and the SID sub-TLVs in which each type MAY Appear, are consistent with that in [\[I-D.ietf-lsr-isis-srv6-extensions\]](#).

10. Security Considerations

Security concerns for IS-IS are addressed in [ISO10589], [\[RFC5304\]](#), and [\[RFC5310\]](#).

11. Acknowledgements

TBD

12. IANA Considerations

This document requests allocation for the following TLVs and sub-TLVs in the ISIS TLV registry.

12.1. SRv6 Locator per TN-slice TLV

This document adds one new TLV to the IS-IS TLV Codepoints registry.

Value: TBA1

Name: SRv6 Locator for specific TN-slice

12.2. SRv6 End.X SID per TN-slice and SRv6 LAN End.X SID per TN-slice sub-TLVs

This document adds the definition of two new sub-TLVs in the "sub-TLVs for TLV 22, 23, 25, 141, 222 and 223 registry".

Type: TBA2

Description: SRv6 End.X SID per TN-slice

Type: TBA3

Description: SRv6 LAN End.X SID per TN-slice

13. Normative References

[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-04](#) (work in progress), January 2020.

[I-D.ietf-spring-srv6-network-programming]

Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-ietf-spring-srv6-network-programming-08](#) (work in progress), January 2020.

[I-D.peng-teas-network-slicing]

Peng, S., Chen, R., Mirsky, G., and F. Qin, "Packet Network Slicing using Segment Routing", [draft-peng-teas-network-slicing-02](#) (work in progress), December 2019.

[I-D.zch-lsr-isis-network-slicing]

Zhu, Y., Chen, R., Peng, S., and F. Qin, "IS-IS Extensions to Support Packet Network Slicing using Segment Routing", [draft-zch-lsr-isis-network-slicing-03](#) (work in progress), December 2019.

[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>>.

[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>>.

Authors' Addresses

Shaofu Peng
ZTE Corporation

Email: peng.shaofu@zte.com.cn

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn

Gregory Mirsky
ZTE Corporation

Email: gregimirsky@gmail.com

