LSR WG
Internet-Draft
Intended status: Standards Track
Expires: November 9, 2020

Ran. Chen
Shaofu. Peng
ZTE Corporation
May 8, 2020

IGP Extensions for Shorter SRv6 SID
draft-chen-lsr-igp-shorter-srv6-extensions-02

Abstract

This document describes the IGP extensions required to support the Shorter SRv6 SIDs(Compressing SRv6 SIDs).

Table of Contents

1. Introduction

Segment Routing [RFC8402] leverages the source routing paradigm. An ingress node steers a packet through an ordered list of instructions,called segments.

Segment Routing can be directly instantiated on the IPv6 data plane through the use of the Segment Routing Header defined in [RFC8754]. SRv6 refers to this SR instantiation on the IPv6 dataplane.

However, the size of the SRv6 SID presents a scalabilities challenge to use topological instructions that define a strict explicitly routed path in combination with service-based instructions. At the same time, the size of the SRH/SID may be a challenge for some data plane processors and traffic overhead. [I-D.cheng-spring-shorter-srv6-sid-requirement] describes a list of requirements for the use of a shortened identifier in a segment routing network with the IPv6 data plane.

[I-D.mirsky-6man-unified-id-sr] proposed an extension of SRH that enables the use of a shorter segment identifier in dataplane, such as 32-bits Label format SID or 32-bits IP address format SID.

This document defines extensions to IGP in order to to support the Shorter SRv6 SIDs contained in SID list that installed in dataplane.

2. Advertising Shorter SRv6 SIDs capabilities.

2.1. IS-IS Extensions

A node indicates that it supports the SR Segment Endpoint Node functionality as specified in [RFC8754] by advertising a new SRv6 Capabilities sub-TLV [I-D.ietf-lsr-isis-srv6-extensions] of the router capabilities TLV [RFC7981].

This document extensions the flags field in the SRv6 Capabilities sub-TLV [I-D.ietf-lsr-isis-srv6-extensions] to indicate the node supports the Shorter SRv6 SIDs.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Type        |     Length    |             Flags           |U|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Optional sub-sub-TLVs...
```

Figure 1: U-Flag in SRv6 Capabilities sub-TLV

where

U: U-SID Encapsulation Capability. When the U flag is set, it indicates that the node supports the encapsulate and decapsulate the U-SID, that is to say, the SID list composed of multiple classic 128 bit SIDs can be compressed into an U-SID list containing multiple shorter U-SIDs, which is encapsulated in SRH, or the shorter U-SID can be obtained from SRH and restored to the classic 128 bit SID.

Optional sub-sub-TLVs: When the U flag is set, A new U-Domain sub-sub-TLV is carried to describe which compression domain (U-Domain) the node is in. If the U-Domain sub-sub-tlv is not carried, it is in 32-bit compression domain by default. Note that each node is always in the classical 128 bit compression domain, without explicit notification.

The format of the U-Domain sub-Sub-TLV is as below:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| sub-Type=U-D  |     Length    |             Flags       |M|S|T|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2:U-Domain sub-Sub-TLV

Three flags are currently defined in the U-Domain sub-Sub-TLV:

T: The node is in 32-bit compression domain.

S: The node is in 16-bit compression domain.

M: The node is in MPLS compression domain.

Note that the U-SID Encapsulation capability has nothing to do with the type of compression domain the node is in. For example, an N1 node in a 128 bit compression domain has U-SID Encapsulation capability, while an N2 node in the same domain may not have U-SID Encapsulation capability.

2.2. OSPFv3 Extensions

The SRv6 Capabilities TLV is used by an OSPFv3 router to advertise its SRv6 support along with its related capabilities for SRv6 functionality. This is an optional top level TLV of the OSPFv3 Router Information LSA [RFC7770] which MUST be advertised by an SRv6 enabled router.

This document extensions the flags field in the SRv6 Capabilities TLV [I-D.ietf-lsr-ospfv3-srv6-extensions] to indicate the node supports the Shorter SRv6 SIDs.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Type            |          Length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Flags           |U|            Reserved             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Sub-TLVs...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: U-Flag in SRv6 Capabilities TLV

where

U: U-SID Encapsulation Capability. When the U flag is set, it indicates that the node supports the encapsulate and decapsulate the U-SID, that is to say, the SID list composed of multiple classic 128 bit SIDs can be compressed into an U-SID list containing multiple shorter U-SIDs, which is encapsulated in SRH, or the shorter U-SID can be obtained from SRH and restored to the classic 128 bit SID.

Sub-TLVs: When the U flag is set, A new U-Domain sub-TLV is carried to describe which compression domain (U-Domain) the node is in. If the U-Domain sub-tlv is not carried, it is in 32-bit compression domain by default. Note that each node is always in the classical 128 bit compression domain, without explicit notification.

The format of the U-Domain sub-Sub-TLV is as below:

```
        0                   1                   2                   3
        0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
       | sub-Type=U-D  |     Length    |            Flags        |M|S|T|
       +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: U-Domain sub-TLV

Three flags are currently defined in the U-Domain sub-TLV:

T: The node is in 32-bit compression domain.

S: The node is in 16-bit compression domain.

M: The node is in MPLS compression domain.

Note that the U-SID Encapsulation capability has nothing to do with the type of compression domain the node is in. For example, an N1 node in a 128 bit compression domain has U-SID Encapsulation capability, while an N2 node in the same domain may not have U-SID Encapsulation capability.

3. Advertising SRv6 SID Structure Sub-Sub-TLV

SRv6 SID Structure Sub-Sub-TLV is an optional Sub-Sub-TLV of SRv6 End SID Sub-TLV, SRv6 End.U SID Sub-TLV ,and SRv6 LAN End.U SID Sub-TLV .

As discussed in [I-D.ietf-spring-srv6-network-programming], the node with the SRv6 capability will maintain its local SID table. A Local SID is generally composed of two parts, that is, LOC:FUNCT, or may carry arguments at the same time, that is, LOC:FUNCT:ARGS. The controller plane protocol can also use B:N to represent an LOC, where B is SRv6 SID Locator Block and N to represent node N. In other words, the structure of a complete SID is B:N:FUNCT:ARGS.

SRv6 SID Structure Sub-Sub-TLV [I-D.ietf-lsr-isis-srv6-extensions] or SRv6 SID Structure Sub-TLV [I-D.ietf-lsr-ospfv3-srv6-extensions] is used to advertise the length of each individual part of the SRv6 SID.

If a node advertised the compression domains which the node is in, it SHOULD advertise the related SIDs with structure information, otherwise the result optimized SID list will have to contain related classical 128-bits SRv6 SID.

4. Advertising Endpoint Behaviors with U-Flavor

Endpoint behaviors are defined in [I-D.ietf-spring-srv6-network-programming] and[I-D.ietf-6man-spring-srv6-oam] . The codepoints for the Endpoint behaviors are defined in the "SRv6 Endpoint Behaviors" registry defined in [I-D.ietf-spring-srv6-network-programming]. For End, End.X and End.T behaviors, they can also have PSP, USP and USD variants. This document continues to extend the following new flavors for End and End.X behaviors:

U32-Flavor: indicate the next SID is 32-bits IP address.

U16-Flavor: indicate the next SID is 16-bits IP address.

We can take regard the traditional behaviors that has not any indication of next SID type as behaviors with U128-flavor.

To extend the above U related flavors for other endpoint behaviors, such as VPN related SID and SFC related SID, is out the scope of this document.

Note that a SID MUST NOT set two or more of the above flavors at the same time, because these flavors is used to indicate the next SID type in SRH, that is, the local SID entry must provide exact indication for this purpose.

Each of the above U related flavors can be used combined with existing PSP/USP/USD flavors.

5. Operations

Based on the IGP link-state database which contains U-SID Encapsulation Capabilities and SID(s) per U-Flavors, a headend or controller can firstly check which compression domains a computed SR path crossed, then secondly select U-Flavor related SID to construct an optimized E2E SID list.

The detailed description can refer to [I-D.mirsky-6man-unified-id-sr] and [I-D.liu-idr-segment-routing-te-policy-complement].

6. Security Considerations

Procedures and protocol extensions defined in this document do not affect the security considerations discussed in [I-D.ietf-lsr-isis-sr v6-extensions]and[I-D.ietf-lsr-ospfv3-srv6-extensions] .

7. IANA Considerations

TBD

8. Normative References

[I-D.cheng-spring-shorter-srv6-sid-requirement]
Cheng, W., Xie, C., Pang, R., Li, Z., Chen, R., Lijun, L., Duan, X., and G. Mirsky, "Shorter SRv6 SID Requirements", draft-cheng-spring-shorter-srv6-sid-requirement-01 (work in progress), March 2020.

[I-D.ietf-6man-spring-srv6-oam]
Ali, Z., Filsfils, C., Matsushima, S., Voyer, D., and M. Chen, "Operations, Administration, and Maintenance (OAM) in Segment Routing Networks with IPv6 Data plane (SRv6)", draft-ietf-6man-spring-srv6-oam-04 (work in progress), March 2020.

[I-D.ietf-lsr-isis-srv6-extensions]
Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extension to Support Segment Routing over IPv6 Dataplane", draft-ietf-lsr-isis-srv6-extensions-08 (work in progress), April 2020.

[I-D.ietf-lsr-ospfv3-srv6-extensions]
Li, Z., Hu, Z., Cheng, D., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", draft-ietf-lsr-ospfv3-srv6-extensions-00 (work in progress), February 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-15 (work in progress), March 2020.

[I-D.liu-idr-segment-routing-te-policy-complement]
Yao, L. and S. Peng, "BGP Extensions for Unified SID in TE Policy", draft-liu-idr-segment-routing-te-policy-complement-02 (work in progress), May 2020.

[I-D.mirsky-6man-unified-id-sr]
Cheng, W., Mirsky, G., Peng, S., Aihua, L., Wan, X., and C. Wei, "Unified Identifier in IPv6 Segment Routing Networks", draft-mirsky-6man-unified-id-sr-06 (work in progress), March 2020.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 7770, DOI 10.17487/RFC7770, February 2016, <https://www.rfc-editor.org/info/rfc7770>.

[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>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

Authors' Addresses

Ran Chen
ZTE Corporation
No. 50 Software Ave, Yuhuatai Distinct
Nanjing
China

Email: chen.ran@zte.com.cn

Peng Shaofu
ZTE Corporation
No. 50 Software Ave, Yuhuatai Distinct
Nanjing
China

Email: peng.shaofu@zte.com.cn