LSR WG
Internet-Draft
Intended status: Standards Track
Expires: September 29, 2020

Ran. Chen
Shaofu. Peng
ZTE Corporation

March 28, 2020

IGP Extensions for Shorter SRv6 SID
draft-chen-lsr-igp-shorter-srv6-extensions-00

Abstract

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

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 September 29, 2020.

Table of Contents

1. Introduction 2
2. Advertising Shorter SRv6 SIDs capabilities. 2
2.1. IS-IS Extensions 2
2.2. OSPFv3 Extensions 4
3. Advertising SRv6 SID Structure Sub-Sub-TLV 5
4. Advertising Endpoint Behaviors with UET-Flavor 6
5. Operations 7
6. Security Considerations 7
7. IANA Considerations 7
8. Normative References 7
Authors' Addresses 9

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 [I-D.ietf-6man-segment-routing-header]. 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 [I-D.ietf-6man-segment-routing-header] 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.

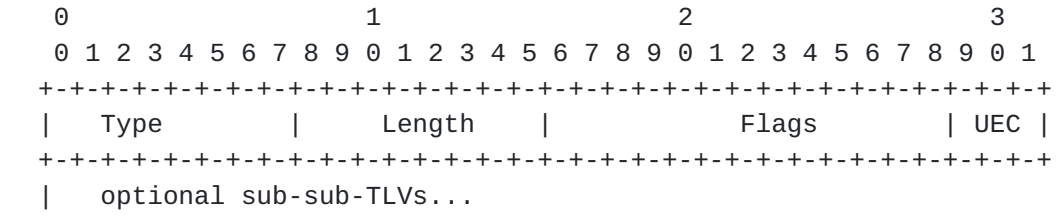

Figure 1: UET-Flag in SRv6 Capabilities sub-TLV

where

UEC: Unified-SID Encapsulation Capability, 3-bits field, refers to [I-D.mirsky-6man-unified-id-sr], it indicates the U-SID capabilities which the node support. A node advertised a specific UEC also means the node belongs to the related UET domain, so it will have capability to install a local SID entry with behavor to get next UET related U-SID from SRH. The value of UEC could be:

0b000: The node only support to use classical 128-bits SRv6 SID. It only belongs to UET-128 domain, and has capability only to get next classical 128-bits SID from SRH.

0b001: The node support to use both classical 128-bits SRv6 SID and 32-bits U-SID. It can belongs to both UET-128 domain and UET-32 IP domain, and has capability both to get next classical 128-bits SID and 32-bits IP U-SID from SRH.

0b010: The node support to use both classical 128-bits SRv6 SID and 32-bits MPLS U-SID. It can belongs to both UET-128 domain and UET-32 MPLS domain, and has capability both to get next classical 128-bits SID and 32-bits MPLS U-SID from SRH.

0b011: The node support to use both classical 128-bits SRv6 SID, 32-bits IP U-SID, and 32-bits MPLS U-SID. It can belongs to both UET-128 domain, UET-32 IP domain, and UET-32 MPLS domain, and has capability both to get next classical 128-bits SID, 32-bits IP U-SID, and 32-bits MPLS U-SID from SRH.

0b100: The node support to use both classical 128-bits SRv6 SID and 16-bits U-SID. It can belongs to both UET-128 domain and

UET-16 domain, and has capability both to get next classical 128-bits SID and 32-bits U-SID from SRH.

0b101: The node support to use both classical 128-bits SRv6 SID, 32-bits IP U-SID, and 16-bits U-SID. It can belongs to both UET-128 domain, UET-32 IP domain, and UET-16 domain, and has capability both to get next classical 128-bits SID, 32-bits IP U-SID, and 16-bits U-SID from SRH.

others: For later defined.

For typical 32-bits based compression scenario, 0b001 UCE is enough.

Note that UEC has two meanings. The first meaning, indicate which UET domain does the advertised node belongs to, this will help to outline which UET domains the SR path crosses. The second meanning, indicate the advertised node has capability to install a local SID entry with UET related behavor, to get next UET related U-SID from SRH, this will help to select appropriate SID with specific UET related behavor for an segment list during compression.

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.

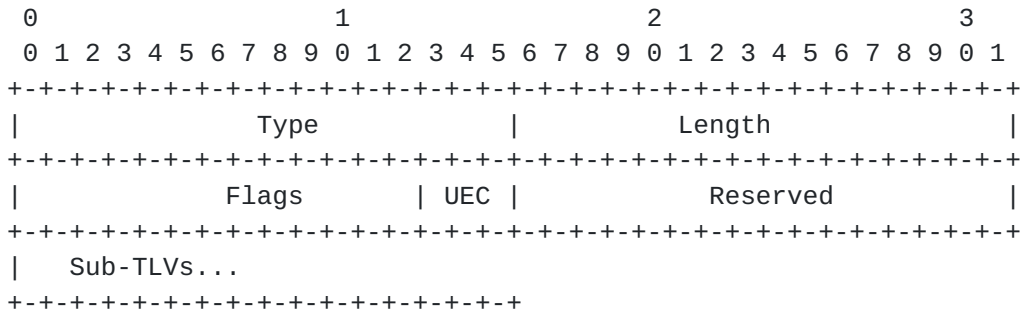

Figure 2: UET-Flag in SRv6 Capabilities TLV

where

UEC: Unified-SID Encapsulation Capability, 3-bits field, refers to [I-D.mirsky-6man-unified-id-sr], it indicates the U-SID capabilities which the node support. A node advertised a specific UEC also means the node belongs to the related UET domain, so it will have capability to install a local SID entry with behavor to get next UET related U-SID from SRH. The value of UEC could be:

0b000: The node only support to use classical 128-bits SRv6 SID. It only belongs to UET-128 domain, and has capability only to get next classical 128-bits SID from SRH.

0b001: The node support to use both classical 128-bits SRv6 SID and 32-bits IP U-SID. It can belongs to both UET-128 domain and UET-32 IP domain, and has capability both to get next classical 128-bits SID and 32-bits IP U-SID from SRH.

0b010: The node support to use both classical 128-bits SRv6 SID and 32-bits MPLS U-SID. It can belongs to both UET-128 domain and UET-32 MPLS domain, and has capability both to get next classical 128-bits SID and 32-bits MPLS U-SID from SRH.

0b011: The node support to use both classical 128-bits SRv6 SID, 32-bits IP U-SID, and 32-bits MPLS U-SID. It can belongs to both UET-128 domain, UET-32 IP domain, and UET-32 MPLS domain, and has capability both to get next classical 128-bits SID, 32-bits IP U-SID, and 32-bits MPLS U-SID from SRH.

0b100: The node support to use both classical 128-bits SRv6 SID and 16-bits U-SID. It can belongs to both UET-128 domain and UET-16 domain, and has capability both to get next classical 128-bits SID and 32-bits U-SID from SRH.

0b101: The node support to use both classical 128-bits SRv6 SID, 32-bits IP U-SID, and 16-bits U-SID. It can belongs to both UET-128 domain, UET-32 IP domain, and UET-16 domain, and has capability both to get next classical 128-bits SID, 32-bits IP U-SID, and 16-bits U-SID from SRH.

others: For later defined.

Other considerations is similar with IS-IS section.

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.X SID Sub-TLV ,and SRv6 LAN End.X 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 an UEC with 0b001/0b011/0b100/0b101, 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 UET-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:

UET-32-IP Flavor: indicate the next SID is 32-bits IP address, termed as UET-1 flavor.

UET-32-MPLS Flavor: indicate the next SID is 32-bits MPLS Label, termed as UET-2 flavor.

UET-16-IP Flavor: indicate the next SID is 16-bits IP address, , termed as UET-3 flavor.

Other flavors are for later defined.

We can take regard the traditional behaviors that has not any indication of next SID type as behaviors with UET-128-IPv6 flavor, termed as UET-0 flavor.

To extend the above UET 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 UET related flavors can be used combined with existing PSP/USP/USD flavors.

If a node supports an UEC, it SHOULD also allocate related SIDs for this UEC, otherwise the result optimized SID list will have to contain related classical 128-bits SRv6 SID.

For example, a node X advertised UCE 0b001, it can allocate a classical END SID X1 with endpoint behavior "End (no PSP, no USP)", it can also allocate an END SID X2 with endpoint behavior "End (no PSP, no USP, UET-32-IP)".

5. Operations

Based on the IGP link-state database which contains UEC capabilities and SID(s) per UET related flavors, a headend or controller can firstly check which UET domains a computed SR path crossed, then secondly select UET 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-srv6-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-segment-routing-header]
Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", draft-ietf-6man-segment-routing-header-26 (work in progress), October 2019.

[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-03 (work in progress), December 2019.

[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-07 (work in progress), March 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-01 (work in progress), March 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>.

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