

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 30, 2020

Yao. Liu
Shaofu. Peng
ZTE Corporation
June 28, 2020

SRv6 Compressed Path Recover Method
draft-pl-spring-compr-path-recover-00

Abstract

This document describes a path information recovery method for compressed SRv6 segment lists.

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 December 30, 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.	Overview	2
3.	Segment List Recover TLV	3
4.	Illustration	5
4.1.	Shifting Method	5
4.2.	Stitching Method	7
4.3.	Mapping Method	8
5.	Security Considerations	8
6.	IANA Considerations	8
7.	References	8
7.1.	Normative References	8
7.2.	Informative References	9
	Authors' Addresses	10

[1.](#) Introduction

Segment routing (SR) [[RFC8402](#)] is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node by inserting an ordered list of instructions, called segments. Segment Routing can be deployed on the MPLS data plane by encoding 32-bits SIDs in MPLS label stack [[RFC8660](#)]. It can also be deployed on the IPv6 data plane by encoding a list of 128-bits SRv6 SIDs in IPv6 Segment Routing Extension Header (SRH) [[RFC8754](#)].

Several proposals are introduced to reduce the overhead for both the traffic volume and the network processor in SRv6. The main ideas of them are basically to use a shorter ID associated with the complete 128 bit SID. The SRH does not carry complete path information after compression, however, the transit and egress nodes may need to know the complete SID information along the path in some scenarios such as OAM.

This document describes a path information recovery method for compressed SRv6 segment lists.

[2.](#) Overview

Several proposals are introduced to reduce the overhead for both the traffic volume and the network processor in SRv6. The main ideas of them are basically to use a shorter ID associated with the complete 128 bit SID. The method to establish such association could be mapping, stitching, shifting, or translation.

Existing compression schemes use the 128 bit space as the carrier of compressing SIDs. In the compressed SRH, the meaning of a 128 bit segment may be:

- a) an uncompressed SID
- b) a unified prefix, followed by several short SIDs, and possibly existing end marker or padding, which is generally set to value 0
- c) short SIDs of the same type and length, possibly followed by end marker or padding, which is generally set to value 0

So after receiving the SRv6 packets, the node or application cannot obtain the complete path information according to segment list in SRH any more.

3. Segment List Recover TLV

This document defines an SRH TLV called Segment List Recover TLV, where,




Figure 1: Segment List Recover TLV

List ID: 16 bits, identification of a complete path with local significance.

List Recover Info: variable length. It provides auxiliary information to recover the complete path information from the segment list in the received SRH.




Figure 2: List Recover Info detailed structure

List Recover Info contains one or more Basic Compressed Description Unit, where:

Index: 8-bit index, number of 128 bit unit remaining from the current carrier, which is a bit like segments left([RFC8200](#), [Section 4.4](#)).

Index refers to the start of (consecutive) carrier(s) that contains the compression information. If a segment list contains both compressed carrier and complete SID, the index will skip the complete SID.

Range: 8 bit. It indicates the number of consecutive 128-bit carrier in the same format, starting from the carrier directed by the index.

In this way, if multiple consecutive carriers use the same format, only one Basic Compressed Description Unit needs to be carried.

When range is set to 1, it indicates that the basic unit carries informational of only one 128 bit carrier.

BL: It indicates the prefix length. When set to 0, it means that the 128 bit carrier does not carry prefix information but only contains short SIDs.

ST: short SID type, 4 bits with following values:

0001: 32-bits SID

0010: 16-bits SID

Existing proposals mainly focus on the 16-bits and 32-bits short SIDs, other values may be defined in the future.

count: 4 bits. It indicates the number of short SIDs in the 128 bit carrier.

When the headend of the SR policy sends the SRv6 packet, it can contain the TLV used to describe the above compression details in the SRH. Upon receiving the above SRv6 packets, the transit node or egress node can perform segment list restoration according to the TLV carried in the SRH.

The restored SID list information can be used as an important reference for constructing the reverse path, or used for displaying to the user or for other OAM purposes.

This document provides two path restoration options:

a) The node locally maintains the mapping between the list ID and the path information. Upon receiving the packet, the node identifies the path corresponding to the SRH according to the list ID carried in the SRH TLV. List Recover Info is not needed in this case.

b) According to the information carried in the list recover info, combined with the segment list in the SRH, the node recovers the complete path information.

It should be noticed that option b can't be used in reduced SRH.

4. Illustration

This section describes how to restore a complete path utilizing a Segment List Recover TLV.

The current chapter focuses on option B combined with some representative schemes. Option A will be discussed in the later version.

4.1. Shifting Method

[I-D.filsfils-spring-net-pgm-extension-srv6-usid] extends SRv6 Network Programming with a new type of SRv6 SID behavior.

A uSID carrier is a 128 bit SRv6 SID of format <uSID-Block><Active-uSID><Next-uSID>...<Last-uSID><End-of-Carrier>...<End-of-Carrier>.

So the segment list may be encoded as figure 3 after compression.

It is consist of six 128-bit units. SID-1 is an uncompressed 128-bit SID and is the first segment to be processed in the segment list.

When restoring the path, the node will save the prefix used in the current carrier, and stitch the prefix, each short SID and possible padding into a complete 128-bit address. If BL in the next Basic Compressed Description Unit is not 0, it indicates that there's a new prefix, and the node will replace the previous prefix with the new prefix, and continue the processing. If BL is 0 as in unit 2 showed in figure 6, the previous prefix will still be used.

SRH may also be used carry short SIDs without any prefix information. In this case, when the node recovers segment list, it may obtain the prefix from the source IP of the IPv6 header.

4.3. Mapping Method

In the mapping method, the short SID is no longer part of the complete 128-bit address, unlike shifting or stitching.

Stitching is supported in [[I-D.mirsky-6man-unified-id-sr](#)].

For the compression proposal using the mapping mode, option A can be used to restore the segment list.

The node need to maintains the mapping between the list ID and the path information. Upon receiving the packet, the node identifies the path corresponding to the SRH according to the list ID carried in the SRH TLV. The list ID can be assigned by a centralized controller or by a node locally.

This will be further discussed in a future version of this document.

5. Security Considerations

The security requirements and mechanisms described in [[RFC8402](#)] and [[RFC8754](#)] also apply to this document. This document does not introduce any new security vulnerabilities.

6. IANA Considerations

TBD

7. References

7.1. Normative References

[I-D.filsfils-spring-net-pgm-extension-srv6-usid]

Filsfils, C., Camarillo, P., Cai, D., Voyer, D., Meilik, I., Patel, K., Henderickx, W., Jonnalagadda, P., Melman, D., Liu, Y., and J. Guichard, "Network Programming extension: SRv6 uSID instruction", [draft-filsfils-spring-net-pgm-extension-srv6-usid-07](#) (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.

[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>>.

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", [RFC 8660](#), DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

[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>>.

7.2. Informative References

[I-D.decraene-spring-srv6-vlsid]

Decraene, B., Raszuk, R., Li, Z., and C. Li, "SRv6 vSID: Network Programming extension for variable length SIDs", [draft-decraene-spring-srv6-vlsid-03](#) (work in progress), March 2020.

[I-D.li-spring-compressed-srv6-np]

Li, Z., Li, C., Xie, C., LEE, K., Tian, H., Zhao, F., Guichard, J., Cong, L., and S. Peng, "Compressed SRv6 Network Programming", [draft-li-spring-compressed-srv6-np-02](#) (work in progress), February 2020.

Authors' Addresses

Liu Yao
ZTE Corporation
No. 50 Software Ave, Yuhuatai District
Nanjing
China

Email: liu.yao71@zte.com.cn

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

Email: peng.shaofu@zte.com.cn

