

Workgroup: 6MAN Working Group
Internet-Draft:
draft-xiao-6man-srv6-checksum-00
Published: 18 May 2022
Intended Status: Standards Track
Expires: 19 November 2022
Authors: X. Min Y. Liu C. Xie
 ZTE Corp. ZTE Corp. China Telecom
 SRv6 Upper-Layer Checksum

Abstract

This document provides a unified mechanism that makes the upper-layer checksum computation rule defined in IPv6 Specification applicable, whether SRv6 SIDs or SRv6 compressed SIDs are used.

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 19 November 2022.

Copyright Notice

Copyright (c) 2022 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](#)
- [2. Conventions](#)
 - [2.1. Requirements Language](#)
 - [2.2. Abbreviations](#)
- [3. Unified Mechanism for Upper-Layer Checksum in SRv6](#)
 - [3.1. C-flag in Segment Routing Header](#)
 - [3.2. C-flag Processing](#)
- [4. IANA Considerations](#)
- [5. Security Considerations](#)
- [6. Acknowledgements](#)
- [7. References](#)
 - [7.1. Normative References](#)
 - [7.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

IPv6 Specification [[RFC8200](#)] defines how upper-layer checksum is computed. Specifically, a "pseudo-header" for IPv6 is constructed as a portion of fields included in upper-layer (e.g., TCP, UDP, ICMPv6, OSPF) checksum computation. As defined in Section 8.1 of [[RFC8200](#)], if the IPv6 packet doesn't contain Routing header, the Destination Address used in the pseudo-header will be in the Destination Address field of the IPv6 header; if the IPv6 packet contains a Routing header, the Destination Address used in the pseudo-header is that of the final destination. In the latter case, at the originating node, that address will be in the last element of the Routing header; at the recipient(s), that address will be in the Destination Address field of the IPv6 header. As also defined in Section 8.1 of [[RFC8200](#)], any node implementing zero-checksum mode of UDP tunnel must follow the requirements specified in "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums" [[RFC6936](#)], and it's outside the scope of this document.

Segment Routing over IPv6 (SRv6) [[RFC8754](#)] defines an IPv6 Routing header called Segment Routing Header (SRH). To comply with the upper-layer checksum computation rule defined in [[RFC8200](#)], at the SRv6 ingress node, the last element of the SRH, i.e., the last Segment Identifier (SID), will become the Destination Address used in the pseudo-header for upper-layer checksum computation; at the SRv6 egress node, after SRH processing is finished, the Destination Address in the IPv6 header will become the Destination Address used in the pseudo-header for upper-layer checksum computation. Note that even at the SRv6 egress node, SRH processing may still invoke IPv6 Destination Address substitution.

The C-SID document [[I-D.ietf-spring-srv6-srh-compression](#)] defines SRv6 compressed SIDs, which use 16-bit or 32-bit SRv6 C-SID to substitute 128-bit SRv6 SID. The NEXT-C-SID flavor and REPLACE-C-SID flavor are defined in the C-SID document. In one case of NEXT-C-SID flavor, at the SRv6 ingress node, the IPv6 packet doesn't contain Routing header, more than one C-SIDs are included in IPv6 Destination Address, the upper-layer checksum computation rule defined in [[RFC8200](#)] doesn't apply anymore. In another case of REPLACE-C-SID flavor, at the SRv6 ingress node, the IPv6 packet contains an SRH, the last element of the SRH is not a 128-bit IPv6 address, but a 16-bit or 32-bit C-SID, the upper-layer checksum computation rule defined in [[RFC8200](#)] doesn't apply anymore.

This document provides a unified mechanism that makes the upper-layer checksum computation rule defined in IPv6 Specification applicable, whether SRv6 SIDs or SRv6 compressed SIDs are used.

2. Conventions

2.1. 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.

2.2. Abbreviations

SR: Segment Routing

SRv6: Segment Routing with IPv6 data plane

SID: Segment ID

C-SID: Compressed Segment ID [[I-D.ietf-spring-srv6-srh-compression](#)]

SRH: Segment Routing Header [[RFC8754](#)]

PSP: Penultimate Segment Pop of the SRH [[RFC8986](#)]

TCP: Transmission Control Protocol [[RFC0793](#)]

UDP: User Datagram Protocol [[RFC0768](#)]

ICMPv6: Internet Control Message Protocol for IPv6 [[RFC4443](#)]

OSPF: Open Shortest Path First protocol [[RFC2328](#)]

3. Unified Mechanism for Upper-Layer Checksum in SRv6

This section defines a unified mechanism for upper-layer checksum in SRv6 networks. This mechanism utilizes a new SRH flag and requests all SRv6 nodes along the transport path to act on the new SRH flag.

3.1. C-flag in Segment Routing Header

[RFC8754] describes the Segment Routing Header (SRH) and how SRv6 capable nodes use it. The SRH contains an 8-bit "Flags" field.

This document defines the following bit in the SRH Flags field to carry the C-flag:

```

    0 1 2 3 4 5 6 7
    +--+--+--+--+--+
    |      |C|      |
    +--+--+--+--+--+
```

Where:

C-flag: Checksum flag in the SRH Flags field defined in [RFC8754]. When C-flag is set, the last element of the SRH MUST be set to an IPv6 address of the final destination.

3.2. C-flag Processing

The C-flag in SRH is used as a marking-bit in the SRv6 packets using upper-layer checksum, each segment endpoint would process the C-flag as defined in this document, to make the SRv6 upper-layer checksum computation smooth and complied to [RFC8200].

At the upper-layer checksum originating node, if the IPv6 packet contains an SRH, the SRH C-flag MUST be set and the Segment List[0] MUST be set to a 128-bit IPv6 address of the final destination; if the IPv6 packet doesn't contain an SRH while the Destination Address field contains more than one compressed SID, an SRH MUST be added with C-flag set and Segment List[0] set to a 128-bit IPv6 address of the final destination. When the upper-layer checksum originating node knows more than one IPv6 address of the final destination, e.g., a local interface address of the final destination, a 128-bit SID locally instantiated at the final destination, and an IPv6 address transformed from a 16-bit or 32-bit compressed SID locally instantiated at the final destination, the originating node needs to select one of them as the last element of SRH, how the originating node makes the choice is beyond the scope of this document.

When the penultimate segment of a segment-list is a Penultimate Segment Pop (PSP) SID, the SRH is removed at the penultimate segment and the C-flag is not processed at the ultimate segment. The

penultimate segment as a PSP SID MUST copy Segment List[0] from the SRH to the Destination Address field of the IPv6 header, then the ultimate segment can still compute the upper-layer checksum with correct IPv6 Destination Address even without SRH.

When an SRv6 node receives a packet destined to S and S is a local SID, the line S01 of the pseudo-code associated with the SID S, as defined in Section 4.3.1.1 of [\[RFC8754\]](#), is appended to as follows for the C-flag processing.

```
S01.2. IF C-flag is set and local configuration permits
  C-flag processing {
    If (Segment List[0] is locally instantiated or represents
        a local interface) {
      a. Set Segments Left to 0.
      b. Update IPv6 DA with Segment List[0].
    }
    Else {
      If (IPv6 DA is locally instantiated as a PSP SID) {
        a. Update IPv6 DA with Segment List[0].
        b. Submit the packet to the egress IPv6 FIB lookup for
            transmission to the new destination.
      }
    }
  }
```

Note that the C-flag processing happens before execution of regular processing of the local SID S. Specifically, the line S01.2 of the pseudo-code specified in this document is inserted between line S01 and S02 of the pseudo-code defined in Section 4.3.1.1 of [\[RFC8754\]](#). When the C-flag defined in this document and the O-flag defined in Section 2.1 of [\[I-D.ietf-6man-spring-srv6-oam\]](#) are both set, the C-flag processing happens after O-flag processing. Specifically, the line S01.2 of the pseudo-code specified in this document is inserted between line S01.1 of the pseudo-code defined in Section 2.1.1 of [\[I-D.ietf-6man-spring-srv6-oam\]](#) and line S02 of the pseudo-code defined in Section 4.3.1.1 of [\[RFC8754\]](#).

Also note that if the final destination needs to be reached more than once on the programmed transport path, the SRv6 packets with C-flag set would be terminated at the first time the final destination is reached. If it's deemed necessary for the SRv6 packets with C-flag set to reach the final destination more than once, more judgment conditions may be added to the pseudo-code of C-flag processing.

4. IANA Considerations

In the "Segment Routing Header Flags" registry created for [\[RFC8754\]](#), a new Checksum Flag is requested from IANA as follows:

Bit Position	Symbol	Description	Semantics Definition	Reference
3	C	Checksum Flag	Section 3.1	This Document

Table 1: New SRH Flag

5. Security Considerations

This document does not raise additional security issues beyond those of the specifications referred to in the list of references.

6. Acknowledgements

TBA.

7. References

7.1. Normative References

[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)", Work in Progress, Internet-Draft, draft-ietf-6man-spring-srv6-oam-13, 23 January 2022, <<https://www.ietf.org/archive/id/draft-ietf-6man-spring-srv6-oam-13.txt>>.

[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>>.

[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>>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

[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.ietf-spring-srv6-srh-compression]

Cheng, W., Filsfils, C., Li, Z., Decraene, B., Cai, D., Voyer, D., Clad, F., Zadok, S., Guichard, J. N., Aihua, L., Raszuk, R., and C. Li, "Compressed SRv6 Segment List Encoding in SRH", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-srh-compression-01, 21 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-srv6-srh-compression-01.txt>>.

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <<https://www.rfc-editor.org/info/rfc768>>.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.

[RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", RFC 6936, DOI 10.17487/RFC6936, April 2013, <<https://www.rfc-editor.org/info/rfc6936>>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

Authors' Addresses

Xiao Min
ZTE Corp.
Nanjing
China

Phone: [+86 25 88013062](tel:+86-25-88013062)
Email: xiao.min2@zte.com.cn

Yao Liu
ZTE Corp.
Nanjing

China

Email: liu.yao71@zte.com.cn

Chongfeng Xie

China Telecom

Beijing

China

Email: xiechf@chinatelecom.cn