

Support Postcard-Based Telemetry for SRv6 OAM
draft-song-6man-srv6-pbt-01

Abstract

Applications such as SRv6 TE may require to collect detailed performance data on SR paths. Existing in-situ OAM techniques incur encapsulation and header overhead issues. This document describes a method based on Postcard-based Telemetry with Packet Marking for SRv6 on-path OAM, which avoids the extra overhead for encapsulating telemetry-related instruction and metadata in SRv6 packets.

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.

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 April 16, 2020.

Copyright Notice

Copyright (c) 2019 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.	PBT Triggered by Marking for SRv6	3
2.1.	Data Template	3
2.2.	Postcard Correlation	4
2.3.	Operational Considerations	4
3.	Use Cases	4
4.	Security Considerations	4
5.	IANA Considerations	4
6.	Contributors	5
7.	Acknowledgments	5
8.	References	5
8.1.	Normative References	5
8.2.	Informative References	5
	Author's Address	6

[1.](#) Introduction

The ability to collect the on-path data about SRv6 packets at each segment is important for SRv6 OAM, especially for monitoring the application-aware services. Some SR-TE algorithms need to acquire realtime flow forwarding performance on each path. The In-situ OAM (IOAM) [[I-D.ietf-ippm-ioam-data](#)] trace option can be used for such purpose. However, SRv6's SRH can be large due to the long segment list. The IOAM trace option introduces significant additional overhead to the SRv6 packets with its instruction and data trace. The large header overhead complicates the packet processing and may exceed the forwarding hardware's header processing capability.

The extra IOAM trace option header also brings encapsulation challenges as documented in [[I-D.li-6man-ipv6-sfc-iftit](#)]. Here we only restate a subtle issue about the IOAM scope: if IOAM header is encapsulated as another IPv6 extension header, the juxtaposition of IOAM and SRH makes it ambiguous to determine the scope and coverage of IOAM: it is unclear if the IOAM is applied to the entire forwarding path or just to the segment nodes. In reality, either case can find its application.

Song

Expires April 16, 2020

[Page 2]

The Direct Export (DEX) option of IOAM described in [\[I-D.ioamteam-ippm-ioam-direct-export\]](#) partially relieves the packet overhead pressure by avoiding including trace data in SRv6 packet, but the encapsulation issue remains, so does the aforementioned ambiguity. In this document, we propose to apply the PBT-M scheme from [\[I-D.song-ippm-postcard-based-telemetry\]](#) for on-path SRv6 telemetry, which can help to solve the encapsulation and overhead issues.

2. PBT Triggered by Marking for SRv6

PBT-M requires marking a packet as a trigger to collect on-path data about the packet. The collected data are exported by an independent "postcard" packet. Therefore, there is no new header encapsulation requirement.

Eight flag bits are currently reserved in SRH. One of those bits can be used as the marking flag, as shown in the following figure. If the "T"-bit is set to 1, the segment node which process the SRH needs to export the on-path data about this packet as pre-configured through management interface.

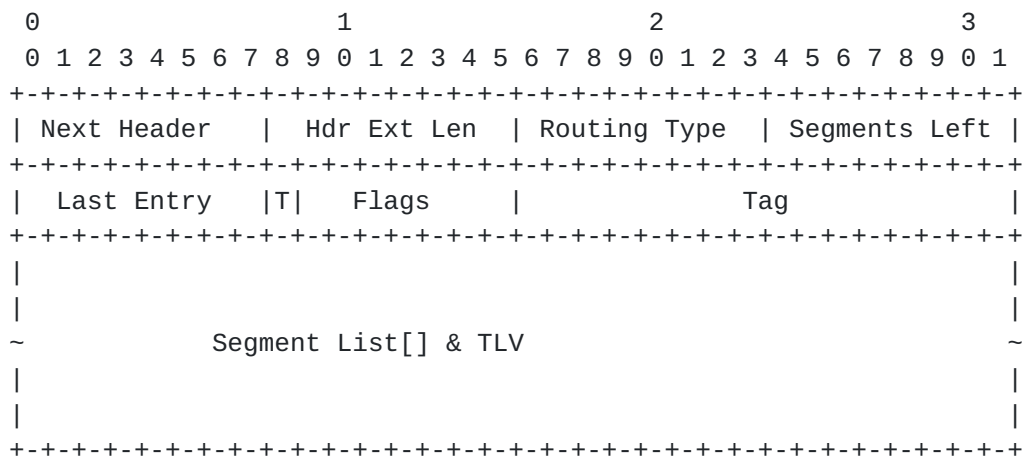


Figure 1: SRH with PBT Mark Flag

2.1. Data Template

It is possible to have the same configuration for all the segment nodes on the data set to collect. However, different flows may require different data collection profiles. It would be more flexible to have multiple different data templates supported by the segment nodes and each packet can designate one template that best

Song

Expires April 16, 2020

[Page 3]

suits its interests to use. The template ID can be carried as a TLV in SRH.

2.2. Postcard Correlation

As discussed in [[I-D.song-ippm-postcard-based-telemetry](#)], PBT-M has some issues to correlate the postcards from the different segment nodes for the same user packet. While several solutions are given to mitigate the problem, it is ideal to be able to correlate the postcards without any constraint and precondition.

A flow ID and a sequence number can be included as TLVs in SRH. The format and usage of the flow ID and the sequence number are the same as those in IOAM DEX option in [[I-D.ioamteam-ippm-ioam-direct-export](#)]. Further, the exported postcard may include the SRH or the current SID which provides a trace to order the postcards.

2.3. Operational Considerations

The SR source node is responsible to determine the policy for setting or resetting the "T"-bit.

A segment node can decide independently whether or not to react on the "T"-bit.

3. Use Cases

TBD.

4. Security Considerations

Since PBT incurs some extra packet processing and transport cost, "T" flag is usually selectively set on a subset of packets by the source node. A potential DoS attack may set the "T" flag for all the packet with the intention to overwhelm the segment nodes. Therefore, the postcards should be generated on the basis of the best effort.

5. IANA Considerations

[I-D.ietf-6man-segment-routing-header] defines a new registry named "Segment Routing Header Flags". This document requests the allocation of a new flag bit "T" for the telemetry trigger mark.

Song

Expires April 16, 2020

[Page 4]

6. Contributors

TBD.

7. Acknowledgments

TBD.

8. References

8.1. Normative References

- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 8126](#), DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.

8.2. Informative References

- [I-D.ietf-6man-segment-routing-header] Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-24](#) (work in progress), October 2019.
- [I-D.ietf-ippm-ioam-data] Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R., daniel.bernier@bell.ca, d., and J. Lemon, "Data Fields for In-situ OAM", [draft-ietf-ippm-ioam-data-07](#) (work in progress), September 2019.
- [I-D.ioamteam-ippm-ioam-direct-export] Song, H., Gafni, B., Zhou, T., Li, Z., Brockners, F., Bhandari, S., Sivakolundu, R., and T. Mizrahi, "In-situ OAM Direct Exporting", [draft-ioamteam-ippm-ioam-direct-export-00](#) (work in progress), October 2019.

Song

Expires April 16, 2020

[Page 5]

[I-D.li-6man-ipv6-sfc-ifit]

Li, Z., Peng, S., and K. LEE, "IPv6 Encapsulation for SFC and IFIT", [draft-li-6man-ipv6-sfc-ifit-02](#) (work in progress), September 2019.

[I-D.song-ippm-postcard-based-telemetry]

Song, H., Zhou, T., Li, Z., Shin, J., and K. Lee, "Postcard-based On-Path Flow Data Telemetry", [draft-song-ippm-postcard-based-telemetry-05](#) (work in progress), September 2019.

Author's Address

Haoyu Song (editor)
Futurewei Technologies
2330 Central Expressway
Santa Clara
USA

Email: hsong@futurewei.com

Song

Expires April 16, 2020

[Page 6]