

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 7, 2022

G. Fioccola
T. Zhou
Huawei
M. Cociglio
Telecom Italia
February 3, 2022

**Segment Routing Header encapsulation for Alternate Marking Method
draft-fz-spring-srv6-alt-mark-02**

Abstract

This document describes how the Alternate Marking Method can be used as the passive performance measurement tool in an SRv6 network. It defines how Alternate Marking data fields are transported as part of the Segment Routing with IPv6 data plane (SRv6) header.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 August 7, 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 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.	Application of the Alternate Marking to SRv6	3
2.1.	Controlled Domain	4
3.	Definition of the SRH AltMark TLV	4
3.1.	Data Fields Format	4
4.	Use of the SRH AltMark TLV	6
5.	Alternate Marking Method Operation	7
6.	Security Considerations	7
7.	IANA Considerations	7
8.	Acknowledgements	8
9.	References	8
9.1.	Normative References	8
9.2.	Informative References	8
	Authors' Addresses	9

[1.](#) Introduction

[RFC8321] and [[RFC8889](#)] describe a passive performance measurement method, which can be used to measure packet loss, latency and jitter on live traffic. Since this method is based on marking consecutive batches of packets, the method is often referred as Alternate Marking Method.

This document defines how the Alternate Marking Method ([[RFC8321](#)]) can be used to measure packet loss and delay metrics for Segment Routing with IPv6 data plane (SRv6).

[RFC8754] defines the Segment Routing Header (SRH) and how it is used by nodes that are Segment Routing (SR) capable.

[I-D.fioccola-v6ops-ipv6-alt-mark] reported a summary on the possible implementation options for the application of the Alternate Marking Method in an IPv6 domain. [[I-D.ietf-6man-ipv6-alt-mark](#)] defines a new TLV that can be encoded in the Option Headers (both Hop-by-hop or Destination) for the purpose of the Alternate Marking Method application in an IPv6 domain.

This document defines how Alternate Marking data is carried as SRH TLV, that can be piggybacked in the packet and transported as part of the SRH. The usage of SRH TLV is introduced in [\[RFC8754\]](#).

2. Application of the Alternate Marking to SRv6

The Alternate Marking Method requires a marking field. A possibility is already offered by [\[I-D.ietf-6man-ipv6-alt-mark\]](#) while the use of a new TLV to be encoded in the SRH is defined in this document.

Since [\[I-D.ietf-6man-ipv6-alt-mark\]](#) defines the IPv6 Application of the Alternate Marking Method through both Hop-by-Hop and Destination Options Header, it is applicable also to SRv6 network. Indeed the use of Destination Option Header carrying Alternate Marking bits coupled with SRH allows to monitor every node along the SR path.

This document introduces the SRH TLV carrying Alternate Marking bits and this can be a preferred approach in case of SRv6 network since it does not rely on the use of Destination Option Header.

The optimization of both implementation and scaling of the Alternate Marking Method is also considered and a way to identify flows is required. The Flow Monitoring Identification field (FlowMonID), as introduced in the next sections, goes in this direction and it is used to identify a monitored flow.

Note that the FlowMonID is different from the Flow Label field of the IPv6 Header ([\[RFC8200\]](#)). Flow Label is used for application service, like load-balancing/equal cost multi-path (LB/ECMP) and QoS. Instead, FlowMonID is only used to identify the monitored flow. The reuse of flow label field for identifying monitored flows is not considered since it may change the application intent and forwarding behaviour. Furthermore the flow label may be changed en route and this may also violate the measurement task. Those reasons make the definition of the FlowMonID necessary for IPv6. Flow Label and FlowMonID within the same packet have different scope, identify different flows, and associate different uses.

An important point that will also be discussed in this document is the uniqueness of the FlowMonID and how to allow disambiguation of the FlowMonID in case of collision.

The following section highlights an important requirement for the application of the Alternate Marking to IPv6 and SRv6. The concept of the controlled domain is explained and it is considered an essential precondition.

2.1. Controlled Domain

[RFC8799] introduces the concept of specific limited domain solutions and, in this regard, it is reported the Application of the Alternate Marking Method as an example.

IPv6 has much more flexibility than IPv4 and innovative applications have been proposed, but for a number of reasons, such as the policies, options supported, the style of network management and security requirements, it is suggested to limit some of these applications to a controlled domain. This is also the case of the Alternate Marking application to SRv6 as assumed hereinafter.

Therefore, the application of the Alternate Marking Method to SRv6 MUST NOT be deployed outside a controlled domain. It is RECOMMENDED that an implementation can be able to reject packets that carry Alternate Marking data and are entering or leaving the controlled domains.

3. Definition of the SRH AltMark TLV

A new TLV carrying the data fields dedicated to the alternate marking method can be defined for the SRH extension headers.

This enables the Alternate Marking Method to take advantage of the network programmability capability of SRv6 ([[I-D.ietf-spring-srv6-network-programming](#)]). Specifically, the ability for an SRv6 endpoint to determine whether to process or ignore some specific SRH TLVs is based on the SID function. The nodes that are not capable of supporting the Alternate Marking functionality do not have to look or process the SRH AltMark TLV and can simply ignore it. This also enables collection of Alternate Marking data only from the supporting segment endpoints.

3.1. Data Fields Format

The following figure shows the data fields format for enhanced alternate marking TLV. This AltMark data is expected to be encapsulated as SRH TLV.

```

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
                                     +---+---+---+---+---+---+---+---+
                                     | SRH TLV Type | SRH TLV Len |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               FlowMonID                               |L|D|   Reserved   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


where:

- o SRH TLV Type: 8 bit identifier of the type of Option/TLV that needs to be allocated. Unrecognised Types MUST be ignored on receipt.
- o SRH TLV Len: The length of the Data Fields of this TLV in bytes.
- o FlowMonID: 20 bits unsigned integer. The FlowMon identifier is described hereinafter.
- o L: Loss flag as defined in [[RFC8321](#)] and [[I-D.ietf-6man-ipv6-alt-mark](#)];
- o D: Delay flag as defined in [[RFC8321](#)] and [[I-D.ietf-6man-ipv6-alt-mark](#)];
- o Reserved: is reserved for future use. These bits MUST be set to zero on transmission and ignored on receipt.

The Flow Monitoring Identification (FlowMonID) is required for some general reasons:

First, it helps to reduce the per node configuration. Otherwise, each node needs to configure an access-control list (ACL) for each of the monitored flows. Moreover, using a flow identifier allows a flexible granularity for the flow definition.

Second, it simplifies the counters handling. Hardware processing of flow tuples (and ACL matching) is challenging and often incurs into performance issues, especially in tunnel interfaces.

Third, it eases the data export encapsulation and correlation for the collectors.

The FlowMon identifier field is to uniquely identify a monitored flow within the measurement domain. The field is set at the source node. The FlowMonID can be uniformly assigned by the central controller or algorithmically generated by the source node. The latter approach cannot guarantee the uniqueness of FlowMonID but it may be preferred for local or private network, where the conflict probability is small due to the large FlowMonID space.

It is important to note that if the 20 bit FlowMonID is set independently and pseudo randomly there is a chance of collision. So, in some cases, FlowMonID could not be sufficient for uniqueness.

This issue is more visible when the FlowMonID is pseudo randomly generated by the source node and there needs to tag it with additional flow information to allow disambiguation. While, in case of a centralized controller, the controller should set FlowMonID by considering these aspects and instruct the nodes properly in order to guarantee its uniqueness.

4. Use of the SRH AltMark TLV

SRv6 leverages the Segment Routing header which consists of a new type of routing header. Like any other use case of IPv6, Hop-by-Hop and Destination Options are useable when SRv6 header is present. Because SRv6 is a routing header, destination options before the routing header are processed by each destination in the route list.

SRH TLV can also be used to encode the AltMark Data Fields for SRv6 and to monitor every node along the SR path. For SRv6, it may be preferred to use the SRH TLV, while for all the other cases with IPv6 data plane the use of the Hop-by-Hop and Destination Option to carry AltMark data fields (as described in [[I-D.ietf-6man-ipv6-alt-mark](#)]) is the best choice.

It is to be noted that the SR nodes implementing the Alternate Marking functionality follows the MTU and other considerations outlined in [[I-D.voyer-6man-extension-header-insertion](#)]. Furthermore, in a SRv6 network, the intermediated nodes that are not in the SID list do not consider the SRH, therefore they cannot support and dig into the SRH TLV.

It is possible to summarize the procedure for AltMark data encapsulation in SRv6 SRH:

- * Ingress Node: As part of the SRH encapsulation, the ingress node of an SR domain or an SR Policy [[I-D.ietf-spring-segment-routing-policy](#)] MAY add the AltMark TLV in the SRH of the data packet, if it supports AltMark functionality and based on local configuration.
- * Intermediate SR Node: The intermediate SR node is any node receiving an IPv6 packet where the destination address of that packet is a local SID. If an intermediate SR node is not capable of processing AltMark TLV, it simply ignores it. While, if an intermediate SR node is capable of processing AltMark TLV, it checks if SRH AltMark TLV is present in the packet using procedures defined in [[RFC8754](#)] and process it.
- * Egress Node: The Egress node is the last node in the segment-list of the SRH. The processing of AltMark TLV at the Egress node

is similar to the processing of AltMark TLV at the Intermediate SR Nodes.

5. Alternate Marking Method Operation

[RFC8321], [RFC8889] describe the Alternate Marking Method in general. While [I-D.ietf-6man-ipv6-alt-mark] describe in detail the application and the Operation of the methodology for IPv6.

6. Security Considerations

The security considerations of SRv6 are discussed in [RFC8754] and [I-D.ietf-spring-srv6-network-programming], and the security considerations of Alternate Marking in general and its application to IPv6 are discussed in [RFC8321] and [I-D.ietf-6man-ipv6-alt-mark].

The Alternate Marking application to IPv6, defined in [I-D.ietf-6man-ipv6-alt-mark], analyzes different security concerns and related solutions. These aspects are valid and applicable also to this document. In particular the fundamental security requirement is that Alternate Marking MUST be applied in a specific limited domain, as also mentioned in [RFC8799].

Alternate Marking is a feature applied to a trusted domain, where one or several operators decide on leveraging and configuring Alternate Marking according to their needs. Additionally, operators need to properly secure the Alternate Marking domain to avoid malicious configuration and attacks, which could include injecting malicious packets into a domain. So the implementation of Alternate Marking is applied within a controlled domain where the network nodes are locally administered. A limited administrative domain provides the network administrator with the means to select, monitor and control the access to the network.

7. IANA Considerations

The SRH TLV Type should be assigned in IANA's Segment Routing Header TLVs Registry.

This draft requests to allocate a SRH TLV Type for Alternate Marking TLV data fields under registry name "Segment Routing Header TLVs" requested by [RFC8754].

SRH TLV Type	Description	Reference

TBD	AltMark Data Fields TLV	This document

8. Acknowledgements

TBD

9. References

9.1. 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>>.

9.2. Informative References

- [I-D.fioccola-v6ops-ipv6-alt-mark]
Fioccola, G., Velde, G. V. D., Cociglio, M., and P. Muley, "IPv6 Performance Measurement with Alternate Marking Method", [draft-fioccola-v6ops-ipv6-alt-mark-01](#) (work in progress), June 2018.
- [I-D.ietf-6man-ipv6-alt-mark]
Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate Marking Method", [draft-ietf-6man-ipv6-alt-mark-12](#) (work in progress), October 2021.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-16](#) (work in progress), January 2022.
- [I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Garvia, P. C., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", [draft-ietf-spring-srv6-network-programming-28](#) (work in progress), December 2020.
- [I-D.voyer-6man-extension-header-insertion]
Voyer, D., Filsfils, C., Dukes, D., Matsushima, S., Leddy, J., Li, Z., and J. Guichard, "Deployments With Insertion of IPv6 Segment Routing Headers", [draft-voyer-6man-extension-header-insertion-10](#) (work in progress), November 2020.

- [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>>.
- [RFC8321] Fioccola, G., Ed., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate-Marking Method for Passive and Hybrid Performance Monitoring", [RFC 8321](#), DOI 10.17487/RFC8321, January 2018, <<https://www.rfc-editor.org/info/rfc8321>>.
- [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>>.
- [RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", [RFC 8799](#), DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.
- [RFC8889] Fioccola, G., Ed., Cociglio, M., Sapio, A., and R. Sisto, "Multipoint Alternate-Marking Method for Passive and Hybrid Performance Monitoring", [RFC 8889](#), DOI 10.17487/RFC8889, August 2020, <<https://www.rfc-editor.org/info/rfc8889>>.

Authors' Addresses

Giuseppe Fioccola
Huawei
Riesstrasse, 25
Munich 80992
Germany

Email: giuseppe.fioccola@huawei.com

Tianran Zhou
Huawei
156 Beiqing Rd.
Beijing 100095
China

Email: zhoutianran@huawei.com

Mauro Cociglio
Telecom Italia
Via Reiss Romoli, 274
Torino 10148
Italy

Email: mauro.cociglio@telecomitalia.it