

6MAN Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 31, 2020

G. Fioccola
T. Zhou
Huawei
M. Cociglio
Telecom Italia
F. Qin
China Mobile
January 28, 2020

IPv6 Application of the Alternate Marking Method
draft-fz-6man-ipv6-alt-mark-05

Abstract

This document describes how the Alternate Marking Method can be used as the passive performance measurement tool in an IPv6 domain and reports implementation considerations. It proposes how to define a new Extension Header Option to encode alternate marking technique and also considers the Segment Routing Header TLV alternative.

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 July 31, 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.	IPv6 application of the Alternate Marking	3
3.	Definition of the AltMark TLV	4
3.1.	Data Fields Format	4
4.	AltMark: EH Option or SRH TLV	5
5.	Alternate Marking Method Operation	6
6.	Security Considerations	6
7.	IANA Considerations	6
8.	Acknowledgements	6
9.	References	7
9.1.	Normative References	7
9.2.	Informative References	7
	Authors' Addresses	8

1. Introduction

[RFC8321] and [[I-D.ietf-ippm-multipoint-alt-mark](#)] 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.

[I-D.song-opsawg-ifit-framework] introduces the telemetry architecture that can be considered as reference.

This document defines how the Alternate Marking Method ([[RFC8321](#)]) can be used to measure packet loss and delay metrics in IPv6.

The format of the IPv6 addresses is defined in [[RFC4291](#)] while [[RFC8200](#)] defines the IPv6 Header, including a 20-bit Flow Label and

the IPv6 Extension Headers. The Segment Routing Header (SRH) is defined in [[I-D.ietf-6man-segment-routing-header](#)].

[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. This document, starting from the outcome of [[I-D.fioccola-v6ops-ipv6-alt-mark](#)], introduces a new TLV that can be encoded in the Option Headers (both Hop-by-hop or Destination) and in the SRH ([[I-D.ietf-6man-segment-routing-header](#)] for the purpose of the Alternate Marking Method application in an IPv6 domain).

2. IPv6 application of the Alternate Marking

The Alternate Marking Method requires a marking field. As mentioned, several alternatives have been analysed in [[I-D.fioccola-v6ops-ipv6-alt-mark](#)] such as IPv6 Extension Headers, IPv6 Address and Flow Label.

The preferred choice would be the use of a new TLV to be encoded in the Option (Hop-by-hop or Destination) header and in the SRH.

This approach is compliant with [[RFC8200](#)] that recommends the use of existing EH rather than defining new ones especially with hop by hop behaviour.

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

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.

Note that the FlowMonID is different from the Flow Label field of the IPv6 Header ([[RFC8200](#)]). Flow Label is used for application service,

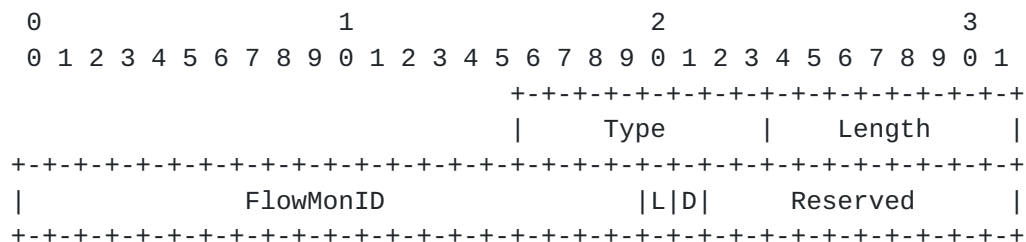
like load-balancing/equal cost multi-path (LB/ECMP) and QoS. Instead, FlowMonID is only used to identify the monitored flow. The reuse of flow label field for identifying monitored flows is not considered since it may change the application intent and forwarding behaviour. Furthermore the flow label 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.

3. Definition of the AltMark TLV

The desired choice is to define a new TLV for the Option and SRH extension headers, carrying the data fields dedicated to the alternate marking method.

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 in the IPv6 Option (hop-by-hop or destination) and SRH extension headers.



where:

- 0 Type/Option Type: 8 bit identifier of the type of Option/TLV that needs to be allocated. Unrecognised Types MUST be ignored on receipt.
- 0 Length/Opt Data Len: The length of the length Data Fields of this Option/TLV in bytes.
- 0 FlowMonID: 20 bits unsigned integer. The FlowMon identifier field is to uniquely identify a monitored flow within the measurement domain. The field is set at the ingress node. The FlowMonID can be uniformly assigned by the central controller or algorithmically generated by the ingress 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.

- o L: Loss flag as defined in [[RFC8321](#)];
- o D: Delay flag as defined in [[RFC8321](#)];
- o Reserved: is reserved for further use. These bits MUST be set to zero on transmission and ignored on receipt.

4. AltMark: EH Option or SRH TLV

Using a new EH Option assumes that all routers in the domain support this type of headers even if an unrecognized EH Option may be just ignored without impacting the traffic. So, the new AltMark Option Layout seems the best way to implement the Alternate Marking method.

It is important to highlight that the Option Layout can be used both as Destination Option and as Hop-By-Hop Option depending on the Use Cases. In general, it is needed to perform both end-to-end and hop-by-hop measurements, and the alternate marking methodology in [[RFC8321](#)] allows, by definition, both performance measurements.

So, Hop-By-Hop Options Header or Destination Options Header can be used based on the chosen type of performance measurement.

SRv6 leverages the Segment Routing header which consists of a new type of routing header. Like any other use case of IPv6, HBH 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. Furthermore, the intermediated nodes that are not in the SID list may consider the SRH as a green field, therefore they cannot support and bypass or support and dig into the SRH TLV.

In summary, it is possible to list the alternative options:

Destination Option => measurement only by node in Destination Address.

Hop-By-Hop Option => every router on the path with feature enabled.

SRH TLV => every node along the SR path.

Destination Option + SRH => every node along the SR path.

Note that the SRH TLV and Destination Option + SRH can be considered equivalent so in this case it may be preferred to use the SRH.

Both [[RFC7045](#)] and [[RFC8200](#)] do not recommend the introduction of new Hop-by-Hop Options headers because nodes may be configured to ignore, drop or assign to a slow processing path. But, in case of the AltMark data fields described in this document, the new hop-by-hop option is needed for OAM and an intermediate node can read it or not but, this does not affect the packet behavior. The source node is the only one that writes the hop-by-hop option to mark alternately the flow, so, the performance measurement can be done for those nodes configured to read this option, while the others are simply not considered for the metrics. Moreover, in case of SRv6, the use of SRH TLV for every node along the SR path is a good choice to implement hop-by-hop measurements.

In addition to the previous alternatives, for legacy network it is possible to mention a non-conventional application of the SRH TLV and Destination Option for the hop-by-hop usage. [[RFC8200](#)] defines that the nodes along a path examine and process the Hop-by-Hop Options header only if HBH processing is explicitly configured. On the other hand, using the SRH TLV or Destination Option for hop-by-hop action would cause worse performance than Hop-By-Hop. The only motivation for hiding the hop-by-hop options inside of destination options can be for compatibility reasons but in general it is not recommended.

5. Alternate Marking Method Operation

[RFC8321] and [[I-D.ietf-ippm-multipoint-alt-mark](#)] describe in detail the methodology.

6. Security Considerations

tbc

7. IANA Considerations

The option type should be assigned in IANA's "Destination Options and Hop-by-Hop Options" registry. Also, the TLV type should be assigned from Segment Routing Header TLVs Registry.

8. Acknowledgements

The authors would like to thank Bob Hinden, Ole Troan, Tom Herbert, Stefano Previdi, Brian Carpenter for the precious comments and suggestions.

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., 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-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-ippm-multipoint-alt-mark]
Fioccola, G., Cociglio, M., Sapio, A., and R. Sisto, "Multipoint Alternate Marking method for passive and hybrid performance monitoring", [draft-ietf-ippm-multipoint-alt-mark-04](#) (work in progress), January 2020.
- [I-D.song-opsawg-ifit-framework]
Song, H., Qin, F., Chen, H., Jin, J., and J. Shin, "In-situ Flow Information Telemetry", [draft-song-opsawg-ifit-framework-10](#) (work in progress), December 2019.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", [RFC 7045](#), DOI 10.17487/RFC7045, December 2013, <<https://www.rfc-editor.org/info/rfc7045>>.
- [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](https://www.rfc-editor.org/info/rfc8321), DOI 10.17487/RFC8321, January 2018, <<https://www.rfc-editor.org/info/rfc8321>>.

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

Fengwei Qin
China Mobile
32 Xuanwumenxi Ave.
Beijing 100032
China

Email: qinfengwei@chinamobile.com

