V6OPS Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 7, 2018

G. Fioccola
Telecom Italia
G. Van de Velde
Nokia
M. Cociglio
Telecom Italia
P. Muley
Nokia

June 5, 2018

IPv6 Performance Measurement with Alternate Marking Method
draft-fioccola-v6ops-ipv6-alt-mark-01

Abstract

This document describes how the alternate marking method in [RFC8321] can be used as the passive performance measurement method in an IPv6 domain, and will discuss the strengths and the weaknesses of the implementation options available to network operations. It proposes how to extend [RFC7837] to apply alternate marking technique.

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 December 7, 2018.

Copyright Notice

Copyright (c) 2018 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

This document reports a summary on the possible implementation options for the application of the alternate marking method in an IPv6 domain.

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

This document defines how the alternate marking method can be used to measure packet loss and delay metrics of IPv6 tunneled packets or SRv6 policies.

The IPv6 Header Format defined in [RFC8200] introduces the format of IPv6 addresses, the Extension Headers in the base IPv6 Header and the availability of a 20-bit flow label, that can be considered for the application of the Alternate Marking methodology.

2. IPv6 application of Alternate Marking

The application of the alternate marking requires a marking field. The alternatives that can be taken into consideration for the choice of the marking field are the following:

- o Extension Header

- o IPv6 Address

- o Flow Label

2.1. IPv6 Extension Headers as Marking Field

A new type of EH may be a solution space proposal (e.g. [RFC8250] and [RFC7837] give a chance).

A possibility can be to use a Hop-By-Hop(HBH) Extension Header(EH). The assumption is that a HBH EH with an alternate marking measurement option can be defined. The router processing can be optimized to handle this use case.

Using a new EH assumes that ALL routers in the domain support this type of headers, which complicates backward compatibility of the technology. The extension of an existing EH (e.g. [RFC7837]) can overcome this issue.

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Option Type  | Option Length |X|L|E|C| MF|res|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Mark Field (MF) is:

```
 0
 0   1
+-+-+-+-+
| S | D |
+-+-+-+-+
```

Figure 1: ConEx HBH Option Layout with Mark Field

where:

o S - Single mark method;

o D - Double mark method.

The Figure 1 defines a new ConEx HBH (Hop-By-Hop) Option Layout.

This proposal starts from ConEx Destination Option Layout defined in [RFC7837], where the Reserved (res) field is made by four bits that are not used in that specification, in fact they are set to zero by the sender and are ignored by the receiver.

This document aims to introduce the Mark Field (2 bits from 4 bits res field). So the Mark Field (MF) reduces the number of Reserved bits and the Reserved (res) field is now made by 2 bits.

It is important to highlight that the Destination Option Layout is used as Hop-By-Hop Option Layout, since the alternate marking methodology in [RFC8321] allows, by definition, Hop-By-Hop performance measurements.

[I-D.krishnan-conex-ipv6] also tried to introduce a ConEx HBH Options and inspired this proposal.

[I-D.fear-ippm-mpdm] introduces Marking Performance and Diagnostic Metrics (M-PDM) and aims to combine [RFC8250] with [RFC8321], while the extension of [RFC7837], proposed in this document, is optimized to include only marking method without any considerations on how to report and manage, this can be done in-band or out-of-band depending on the case.

2.2. Other Possibilities

This section reports the other possibilities that have been discussed.

2.2.1. IPv6 Addresses as Marking Field

There is an advantage of using destination addresses (DA) to encode the alternate marking method. In addition to identifying a host, a destination address is also and more fundamentally identifying an exit point from the forwarding domain. It indicates where processing for forwarding to the DA stops, and where other processing of the packet is to occur. Using the DA to encode this alternate marking processing means that it is easy to retrofit into existing devices and models. There is no need to replace existing IPv6 forwarding devices, because they already support DA based forwarding.

However using DA for marking seems a lot expensive.

2.2.2. IPv6 Flow Label as Marking Field

Considering the Flow Label, [RFC6294] makes a survey of Proposed Use Cases for the IPv6 Flow Label. The flow label is an immutable field recommended to contain a pseudo-random value, however, often it has the default value of zero. [RFC6436] and [RFC6437] open the door for IPv6 Flow Label to be used in a controlled environment and [RFC6438] describes the use of the IPv6 Flow Label field for load distribution purpose, especially across Equal Cost Multi-Path (ECMP) and/or Link Aggregation Group (LAG) paths. In addition it is possible to mention [I-D.krishnan-6man-header-reserved-bits] that tried to set aside 4 bits from the flow label field for future expansion.

There are few drawbacks to use Flow Label instead of an EH solution or IPv6 Addresses for IPv6 alternate marking, in particular an easier backward compatibility and less bits on the wire. In this way nothing breaks if a transit router does not have the capability of understanding the Flow Label context.

Since the flow-label based load balancing has been defined, the application of the Alternate Marking method to the flow label could be realised with two fundamental assumptions:

- o The original flow-label reconstructed when leaving the controlled domain.
- o The usage of IPv6 tunnels (IPv6inIPv6, IPSec, IPv6 UDP, etc..) or SRv6 policies.

In this case, the controlled domain reflects to the fact that it is a network operator choice that grabs control of packet handling within its own network. In fact, regarding the flow label, four options can be supposed:

1) Just do not do anything with Flow Label (leave it default).

2) Entropy only and NO alternate marking for performance measurements.

3) Alternate marking only and NO usage of entropy.

4) Alternate marking and entropy (in this case the entropy SHOULD be based upon a subset of bits because otherwise paths may be changed when the marking changes).

3. Alternate Marking Method Operation

[RFC8321] describes in detail the methodology, that we briefly illustrate also here.

3.1. Single Mark Measurement

As explained in the [RFC8321], marking can be applied to delineate blocks of packets based either on equal number of packets in a block or based on equal time interval. The latter method offers better control as it allows better account for capabilities of downstream nodes to report statistics related to batches of packets and, at the same time, time resolution that affects defect detection interval.

If the Single Mark measurement used, then the D flag MUST be set to zero on transmit and ignored by monitoring point.

The S flag is used to create alternate flows to measure the packet loss by switching value of the S flag. Delay metrics MAY be calculated with the alternate flow using any of the following methods:

o First/Last Batch Packet Delay calculation: timestamps are collected based on order of arrival so this method is sensitive to packet loss and re-ordering.

o Average Packet Delay calculation: an average delay is calculated by considering the average arrival time of the packets within a single block. This method only provides single metric for the duration of the block and it doesn't give information about the delay distribution.

3.2. Double Mark Measurement

Double Mark method allows more detailed measurement of delays for the monitored flow but it requires more nodal and network resources. If the Double Mark method used, then the S flag MUST be used to create the alternate flow. The D flag MUST be used to mark single packets to measure delay jitter.

The first marking (S flag alternation) is needed for packet loss and also for average delay measurement. The second marking (D flag is put to one) creates a new set of marked packets that are fully identified and dedicated for delay. This method is useful to have not only the average delay but also to know more about the statistic distribution of delay values.

4. Security Considerations

tbc

5. IANA Considerations

tbc

6. Acknowledgements

The authors would like to thank Fred Baker, Ole Troan, Robert Hinden, Suresh Krishnan, Brian Carpenter, Roberta Maglione, Tom Herbert, Mark Smith, Joel Halpern, Fernando Gont, Xiaohu Xu and Joel Jaeggli for their comments and feedbacks.

7. References

7.1. 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>.

7.2. Informative References

[I-D.fear-ippm-mpdm]
Elkins, N., Fioccola, G., and m. mackermann@bcbsm.com, "IPv6 Marking and Performance and Diagnostic Metrics (MPDM)", draft-fear-ippm-mpdm-00 (work in progress), June 2018.

[I-D.krishnan-6man-header-reserved-bits]
Krishnan, S. and J. Halpern, "Reserving bits in the IPv6 header for future use", draft-krishnan-6man-header-reserved-bits-00 (work in progress), October 2010.

[I-D.krishnan-conex-ipv6]
Krishnan, S., Kuehlewind, M., and C. Ucendo, "Options for Conex marking in IPv6 packets", draft-krishnan-conex-ipv6-02 (work in progress), March 2011.

[RFC6294] Hu, Q. and B. Carpenter, "Survey of Proposed Use Cases for the IPv6 Flow Label", RFC 6294, DOI 10.17487/RFC6294, June 2011, <https://www.rfc-editor.org/info/rfc6294>.

[RFC6436] Amante, S., Carpenter, B., and S. Jiang, "Rationale for Update to the IPv6 Flow Label Specification", RFC 6436, DOI 10.17487/RFC6436, November 2011, <https://www.rfc-editor.org/info/rfc6436>.

[RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <https://www.rfc-editor.org/info/rfc6437>.

[RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <https://www.rfc-editor.org/info/rfc6438>.

[RFC7837] Krishnan, S., Kuehlewind, M., Briscoe, B., and C. Ralli, "IPv6 Destination Option for Congestion Exposure (ConEx)", RFC 7837, DOI 10.17487/RFC7837, May 2016, <https://www.rfc-editor.org/info/rfc7837>.

[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>.

[RFC8250] Elkins, N., Hamilton, R., and M. Ackermann, "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option", RFC 8250, DOI 10.17487/RFC8250, September 2017, <https://www.rfc-editor.org/info/rfc8250>.

[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>.

Authors' Addresses

Giuseppe Fioccola
Telecom Italia
Torino
Italy

Email: giuseppe.fioccola@telecomitalia.it

Gunter Van de Velde
Nokia
Antwerp
BE

Email: gunter.van_de_velde@nokia.com

Mauro Cociglio
Telecom Italia
Torino
Italy

Email: mauro.cociglio@telecomitalia.it

Praveen Muley
Nokia
Mountain View
USA

Email: praveen.muley@nokia.com