V6OPS Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 30, 2018

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

February 26, 2018

IPv6 Performance Measurement with Alternate Marking Method
draft-fioccola-v6ops-ipv6-alt-mark-00

Abstract

This document describes how the alternate marking method 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.

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 30, 2018.

Table of Contents

1. Introduction

This document reports a summary on the possible implemetation 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] and [RFC2460] introduces the availability of an 20-bit flow label, the format of IPv6 addresses and the Extension Headers in the base IPv6 Header.

For instance, 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. It is important to underline that these specifications encourage non-zero flow label values to be used and clearly defines how to set a non-zero value and it retains the rule that the flow label must not be changed en route but allows routers to set the label on behalf of hosts that do not do so. Based on these considerations, it is allowed to use the flow label field in a managed domain, assuming that, when a packet leaves the domain, the original flow label value MUST be restored or the packet MUST be found invalid.

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:

- Extension Header
- IPv6 Address
- Flow Label

2.1. IPv6 Extension Headers as Marking Field

A new type of EH may be a solution space proposal (e.g. [RFC8250] gives 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 case.

2.2. 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.

2.3. IPv6 Flow Label as Marking Field

There are few other drawbacks to use Flow Label instead of an EH solution or IPv6 Addresses for IPv6 alternate marking:

- easier backward compatibility because nothing breaks if a transit router does not have the capability of understanding the Flow Label context (in that case the flow-label in the outer tunnel header is just a flow-label).

- having a EH seems less backward compatible, and will be less easy to use unless ALL routers in the domain support these type of headers.

- using DA for marking seems expensive.

- For most of the routers the support nearly comes for free.

- Less bits on the wire (going SRv6 has already a significant bits-on-wire tax because of the outer IPv6 header and the SRv6 EH).

So, using the flow-label in the outer IPv6 tunnel header (e.g. SRv6 header) gives some benefits. The flow-label as marking field, is basically something that routers can do right now, and it does not break any IPv6 rules and is expected to be supported by the routers by default. Indeed the solution proposed in the draft, is for the moment assumed to be exclusive from other usages (with exception of entropy) of the flow-label in the controlled operator domain. Currently, network operators traditionally do not use flow-label at all, hence the above assumption seems reasonable.

So, the application of the Alternate Marking method in a managed and controlled domain could be realised with two fundamental assumptions:

- The original flow-label reconstructed when leaving the controlled domain.

- The usage of IPv6 tunnels (IPv6inIPv6, IPSec, IPv6 UDP, etc..) or SRv6 policies.

The Figure 1 displays format of the possible Mark Field (2 bits from 20 bit IPv6 flow-label field).

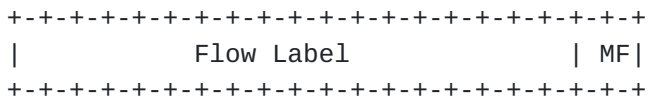

Mark Field (MF) is:

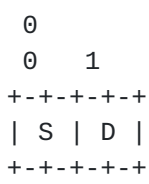

Figure 1: Mark field format

where:

- S - Single mark method;

- D - Double mark method.

The use of the other 18 bits is not specified in this document because is out of scope here. But it should follow [RFC6437], where flow-label based load balancing, ECMP or LAG is described. The methodology SHOULD be used within a controlled domain where the load-balancing based on flow label is disabled. Otherwise, the network elements MUST mask the Mark Field (MF), so it will not change hashing calculation for the same flow because only 18 bits + 2 zeros can be used for the entropy.

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. The network operator adds through policy the outer SRv6 header and has in fact three options regarding flow label:

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

2) Alternate marking only and NO usage of entropy.

3) Alternate marking and entropy (in this case the entropy SHOULD be based upon 18 bits instead of 20 bits because otherwise paths

may be changed when the marking changes (e.g. periods of 5 minutes per marking period). This is however not a MUST because some operators do not care if because of marking change.

The closed system here is defined by the IPv6 Tunnel or SRv6, in particular it is the network between the head-end (where the outer header is added) and the tail-end (where the outer header is removed).

2.3.1. IPv6 Tunnel Use Case

The ingress router is the "source" of the IPv6 tunnel and impose the OUTER IPv6 header, so Ingress router can control 2 bits (Mark Field) from 20 bit flow-label field of OUTER IPv6 header. The Egress router removes OUTER IPv6 header, restoring ORIGINAL payload and payload headers (IPv6, IPv4, L2 traffic, MBH, etc...).

The flow-label is set "only" by the tunnel head-end router on the outer IPv6 header. The tunnel head-end router can do this because it is the device that created the outer header. The original IPv6 packet is riding inside the tunnel, and as result the original flow-label and original IPv6 header is left untouched.

2.3.2. SRv6 Use Case

When IPv6 SRv6 Encapsulation is used, the outer SRv6 header uses 2 bits (Mark Field) from 20 bit flow-label field. Outer SRv6 header will be removed when exiting the SP domain and the original flow-label is restored at egress.

The flow label of the original packet is untouched. The flow label that is set in this proposal is done at the SRv6 tunnel head-end which imposes the SRv6 encapsulation header. So basically, it is just the SRv6 tunnel outer encap header which is used for alternate marking. And this is set only one time by the original SRv6 tunnel head-end router (which is the source address of the IPv6 SRv6 tunnel). This outer SRv6 header is removed when the packet exits the SRv6 domain, and the original flow label appears again untouched. So, in this proposal there is no device which is changing flow-labels at all. It is only during the imposing of the SRv6 outer header, that the flow label field is set once for Alternate marking purposes inside the outer SRv6 tunnel header.

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:

- 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.

- 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. Acknowledgements

The authors would like to thank Brian Carpenter, Fred Baker, Tom Herbert, Mark Smith, Joel Halpern, Fernando Gont, Xiaohu Xu and Joel Jaeggli for their comments and feedbacks.

6. IANA Considerations

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

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <https://www.rfc-editor.org/info/rfc2460>.

[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>.

[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