

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 29, 2018

G. Fioccola, Ed.
Telecom Italia
G. Van de Velde, Ed.
Nokia
M. Cociglio
Telecom Italia
P. Muley
Nokia
October 26, 2017

Using the IPv6 Flow Label for Performance Measurement with Alternate
Marking Method in Segment Routing
draft-fioccola-spring-flow-label-alt-mark-01

Abstract

[RFC6294] makes a survey of Proposed Use Cases for the IPv6 Flow Label. The IPv6 protocol includes a flow label in every packet header, but this field is, according to [\[RFC6294\]](#), not used in practice. This document describes how the alternate marking method can be used as the passive performance measurement method in a IPv6 domain.

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 April 29, 2018.

Internet-Draft

IPv6 Flow Label PM with AMM in SR

October 2017

Copyright Notice

Copyright (c) 2017 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 Flow Labels and Alternate Marking	3
2.1.	IPv6 Tunnel	3
2.2.	SRv6 Policy	4
3.	Alternate Marking Method Operation	4
3.1.	Single Mark Measurement	5
3.2.	Double Mark Measurement	5
4.	Security Considerations	6
5.	Acknowledgements	6
6.	IANA Considerations	6
7.	References	6
7.1.	Normative References	6
7.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Introduction

The IPv6 Header Format defined in [\[RFC2460\]](#) introduces the availability of an 20-bit flow label in the base IPv6 Header.

The flow label is an immutable field recommended to contain a pseudo-random value [\[RFC6437\]](#). However, in most packets it has the default value of zero, although some TCP/IP stacks do set it according to [\[RFC6437\]](#).

[\[RFC6436\]](#) and [\[RFC6437\]](#) open the door for IPv6 Flow Label to be used

in controlled environment, e.g. for IPv6 tunnelled packets or SRv6 policies (see following Sections for more details). For example [[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.

But, a specific goal presented in [[RFC6437](#)] is to enable and encourage the use of the flow label. In fact it is important to underline two aspects from this specification:

- o It encourages non-zero flow label values to be used and clearly defines how to set a non-zero value.
- o 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 when a packet leaves the domain, the original flow label value MUST be restored.

[I-D.ietf-ippm-alt-mark] 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.

[2.](#) IPv6 Flow Labels and Alternate Marking

The application of the Alternate Marking method in a managed and controlled domain is realised with two fundamental assumptions:

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

[2.1.](#) IPv6 Tunnel

The ingress router is the "source" of the IPv6 tunnel and impose the OUTER IPv6 header, so Ingress router can control the Flow-label of OUTER IPv6 header. The Egress router removes OUTER IPv6 header, restoring ORIGINAL payload and payload headers (IPv6, IPv4, L2 traffic, MBH, etc...).

[2.2.](#) SRv6 Policy

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.

As described in [[I-D.voyer-6man-extension-header-insertion](#)], when IPv6 SRv6 EH insertion is used, there is the insertion of IPv6 Extension Headers. However in [[RFC8200](#)] it was ruled against insertion of EHs and it is recommended that Extension Headers are not processed, inserted, or deleted by any node along the path, until the packet from source node reaches the destination node.

[3.](#) Alternate Marking Method Operation

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

```
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Flow Label           | MF |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Mark Field (MF) is:

```
0
0  1
```

```

+--+--+--+
| S | D |
+--+--+--+

```

Figure 1: Mark field format

where:

- o S - Single mark method;
- o 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 [\[RFC6437\]](#). If the domain is using flow-label based load balancing, ECMP or LAG internally, there are two possibilities. The methodology SHOULD be used within a controlled domain where the load-balancing based on flow label is disabled. Otherwise, the network element MUST mask the

Mark Field (MF), so it will not change hashing calculation for the same flow because only 18 bits + 2 zeros are used for the entropy.

[3.1.](#) Single Mark Measurement

As explained in the [\[I-D.ietf-ippm-alt-mark\]](#), 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.](#) Acknowledgements

tbc

[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

- [I-D.ietf-ippm-alt-mark]
Fioccola, G., Capello, A., Cociglio, M., Castaldelli, L.,
Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi,
"Alternate Marking method for passive and hybrid
performance monitoring", [draft-ietf-ippm-alt-mark-13](#) (work
in progress), October 2017.
- [I-D.voyer-6man-extension-header-insertion]
daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d.,
Leddy, J., Filsfils, C., Previdi, S., Salsano, S.,
Cianfrani, A., Lebrun, D., Bonaventure, O., Jonnalagadda,
P., Sharif, M., Elmalky, H., Abdelsalam, A., Raszuk, R.,
Ayyangar, A., Steinberg, D., and W. Henderickx, "Insertion
of IPv6 Segment Routing Headers in a Controlled Domain",
[draft-voyer-6man-extension-header-insertion-01](#) (work in
progress), September 2017.
- [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>>.

Authors' Addresses

Giuseppe Fioccola (editor)
Telecom Italia
Torino
Italy

Email: giuseppe.fioccola@telecomitalia.it

Gunter Van de Velde (editor)
Nokia
Antwerp
BE

Email: gunter.van_de_velde@nokia.com

Mauro Cociglio
Telecom Italia
Torino
Italy

Email: mauro.cociglio@telecomitalia.it

Nokia
Mountain View
USA

Email: praveen.muley@nokia.com