

SFC Working Group
Internet-Draft
Intended status: Experimental
Expires: 28 March 2022

G. Mirsky
Ericsson
G. Fioccola
Huawei Technologies
T. Mizrahi
Huawei Network.IO Innovation Lab
24 September 2021

Performance Measurement (PM) with Alternate Marking Method in Service
Function Chaining (SFC) Domain
draft-mirsky-sfc-pmamm-14

Abstract

This document describes how the alternate marking method can be used as the efficient performance measurement method taking advantage of the actual data flows in a Service Function Chaining (SFC) domain.

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 28 March 2022.

Copyright Notice

Copyright (c) 2021 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. Conventions used in this document	3
2.1. Acronyms	3
2.2. Requirements Language	3
3. Mark Field in NSH Base Header	3
4. Theory of Operation	4
4.1. Single Mark Enabled Measurement	5
4.2. Multiplexed Mark Enabled Measurement	5
4.3. Residence Time Measurement with the Alternate Marking Method	6
5. IANA Considerations	6
5.1. Mark Field in NSH Base Header	6
6. Security Considerations	7
7. Acknowledgment	7
8. References	7
8.1. Normative References	7
8.2. Informative References	7
Authors' Addresses	8

1. Introduction

[RFC7665] introduced the architecture of a Service Function Chain (SFC) in the network and defined its components. These include Classifier, Service Function Forwarder (SFF), Service Function (SF), and Service Function proxy. [RFC8924] provides a reference framework for Operations, Administration and Maintenance (OAM) for SFC. [RFC8321] describes the hybrid 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 procedure is often referred to as Alternate Marking Method (AMM).

This document defines how packet loss and delay metrics of a service flow over end-to-end (E2E) Service Function Path (SFP) or any SFP segment can be measured using AMM. This document is aligned with the SFC OAM Performance Measurement requirements defined in [RFC8924]. It states that any SFC-aware network device must have the ability to

perform loss and delay measurements over the service function chain as a unit, i.e., E2E, or to a specific segment of service function through the SFC. Besides, AMM can be used in combination with [I-D.ietf-sfc-ioam-nsh] complementing it in achieving the SFC performance measurement objective.

2. Conventions used in this document

2.1. Acronyms

AMM: Alternate Marking Method

OAM: Operations, Administration and Maintenance

SFC: Service Function Chain

SFP: Service Function Path

SF: Service Function

SFF: Service Function Forwarder

SPI: Service Path Identifier

NSH: Network Service Header

E2E end-to-end

2.2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Mark Field in NSH Base Header

[RFC8300] defines the format of the Network Service Header (NSH).




Figure 1: NSH Base format

This document defines the one-bit long field, referred to as Mark field (M in Figure 1, as part of NSH Base and designated for the alternate marking performance measurement method [RFC8321]. The Mark field MUST be set to 0 at initialization of NSH and ignored on the receipt when the method is not in use. The Mark field MUST NOT be used in defining forwarding and/or quality of service treatment of an SFC packet. The Mark field MUST be used only for the performance measurement of data traffic in the SFC layer. Though the setting of the field to any value likely not affect forwarding and/or quality of service treatment of a packet, the alternate marking method in the SFC layer is characterized as an example of a hybrid performance measurement method according to [RFC7799].

4. Theory of Operation

The marking method can be successfully used in the SFC. Without limiting any generality consider SFC presented in Figure 2. Any combination of markings, Loss and/or Delay, can be applied to a service flow by any SFC component at either ingress or egress point to perform node, link, segment, or E2E measurement to detect performance degradation defects and localize them efficiently.




Figure 2: SFC network

An SFP might include a Re-classifier. Processing of an SFC packet by the Re-classifier might result in that packet being directed to a different SFP identified, for example, by Service Path Identifier's (SPI) value A'. In that case, the Re-classifier MUST set the value of the Mark field according to the local AMM policy defined for the SPI's value A'. Note that the default AMM policy is to set the value of the Mark field to 0.

Using the marking method, a component of the SFC creates distinct sub-flows in the particular service traffic over SFC. Each sub-flow consists of consecutive blocks that are unambiguously recognizable by a monitoring point at any component of the SFC and can be measured to calculate packet loss and/or packet delay metrics.

4.1. Single Mark Enabled Measurement

As explained in the [RFC8321], marking can be applied to delineate blocks of packets based either on the equal number of packets in a block or based on the same time interval. The latter method offers better control as it allows a 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.

The Mark flag is used to create distinctive flows to measure the packet loss by switching the value of the Mark flag every N-th packet or at specified time intervals. Delay metrics MAY be calculated with the alternate flow using any of the following methods:

- * First/Last Packet Delay calculation: whenever the marking, i.e., the value of Mark flag changes a component of the SFC can store the timestamp of the first/last packet of the block. The timestamp can be compared with the timestamp of the packet that arrived in the same order through a monitoring point at a downstream component of the SFC to compute packet delay. Because timestamps collected based on the order of arrival, this method is sensitive to packet loss and re-ordering of packets
- * Average Packet Delay calculation: an average delay is calculated by considering the packets' average arrival time within a single block. A component of the SFC may collect timestamps for each packet received within a single block. The timestamp average is the sum of all the timestamps divided by the total number of packets received. Then the difference between averages calculated at two monitoring points is the average packet delay on that segment. This method is robust to out-of-order packets and packet loss (only a small error is introduced). This method only provides a single metric for the duration of the block, and it doesn't give the minimum and maximum delay values. Highly optimized implementation of the method can reduce the duration of the block and thus overcome the limitation.

4.2. Multiplexed Mark Enabled Measurement

There is also a scheme that method allows measurement of minimum and maximum delays for the monitored flow using a single marking flag. This methodology is described in [I-D.mizrahi-ippm-compact-alternate-marking]. The concept is that in the middle of each block of packets with a certain value of the M flag, a single packet has the M flag inverted. By examining the stream, the packets with the inverted bit can be easily identified and employed for delay measurement. This variation of AMM is

advantageous because it requires only one bit from each packet, and such bits are always in short supply.

4.3. Residence Time Measurement with the Alternate Marking Method

Residence time is the variable part of the propagation delay that a packet experiences while traversing a network, e.g., SFC. Residence Time over an SFC is the sum of the nodal residence times, i.e., periods that the packet spent in each SFF that composes the SFC. The nodal residence time in SFC itself is the sum of sub-nodal residence times that the packet spent in each of SFs that are part of the given SFC and are mapped to the SFF. The residence time and deviation of the residence time metrics may include any combination of minimum, maximum values over the measurement period. It also may include mean, median, and percentiles calculated values. These metrics may be used to evaluate the performance of the SFC and its elements before and during its operation.

Use of the specially marked packets simplifies residence time measurement and correlation of the measured metrics over the E2E SFC. For example, AMM may be used as described in Section 4.2 to identify packets in the data flow to be used to measure the residence time. The nodal and sub-nodal residence time metrics can be locally calculated and then collected using either in-band or out-band OAM mechanisms.

5. IANA Considerations

5.1. Mark Field in NSH Base Header

This document requests IANA to allocate the one-bit field from NSH Base Header Bits [RFC8300] as the Mark field of NSH as the following:

Bit Position	Description	Reference
TBA	Mark field	This document

Table 1: Mark field of SFC NSH

6. Security Considerations

This document defines the use of AMM in an SFC domain and thus all security considerations specific to SFC discussed in [RFC7665] and [RFC8300] are applicable. By introducing AMM into the SFC environment, it inherits all security considerations discussed in [RFC8321]. A new Mark flag is defined in this specification to be used by AMM. Processing of AMM does require additional computational resources and creates a certain amount of state information per AMM flow performance metrics. An implementation MUST provide control over the number of concurrent AMM flows that a node process.

7. Acknowledgment

TBD

8. References

8.1. 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>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", RFC 8300, DOI 10.17487/RFC8300, January 2018, <<https://www.rfc-editor.org/info/rfc8300>>.

8.2. Informative References

- [I-D.ietf-sfc-ioam-nsh] Brockners, F. and S. Bhandari, "Network Service Header (NSH) Encapsulation for In-situ OAM (IOAM) Data", Work in Progress, Internet-Draft, draft-ietf-sfc-ioam-nsh-06, 31 July 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-sfc-ioam-nsh-06>>.
- [I-D.mizrahi-ippm-compact-alternate-marking] Mizrahi, T., Arad, C., Fioccola, G., Cociglio, M., Chen, M., Zheng, L., and G. Mirsky, "Compact Alternate Marking Methods for Passive and Hybrid Performance Monitoring", Work in Progress, Internet-Draft, draft-mizrahi-ippm-

compact-alternate-marking-05, 6 July 2019,
<<https://datatracker.ietf.org/doc/html/draft-mizrahi-ippm-compact-alternate-marking-05>>.

- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <<https://www.rfc-editor.org/info/rfc7665>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [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>>.
- [RFC8924] Aldrin, S., Pignataro, C., Ed., Kumar, N., Ed., Krishnan, R., and A. Ghanwani, "Service Function Chaining (SFC) Operations, Administration, and Maintenance (OAM) Framework", RFC 8924, DOI 10.17487/RFC8924, October 2020, <<https://www.rfc-editor.org/info/rfc8924>>.

Authors' Addresses

Greg Mirsky
Ericsson

Email: gregimirsky@gmail.com

Giuseppe Fioccola
Huawei Technologies

Email: giuseppe.fioccola@huawei.com

Tal Mizrahi
Huawei Network.IO Innovation Lab
Israel

Email: tal.mizrahi.phd@gmail.com