

IPPM Working Group
Internet-Draft
Intended status: Standards Track
Expires: 19 April 2026

L. Zhang
T. Zhou
Huawei
J. Ma
CAICT
16 October 2025

In Situ Operations, Administration, and Maintenance (IOAM) Active
Measurement for Multi-path
draft-zhang-ippm-ioam-active-mp-01

Abstract

Active measurements are typically used to collect the information of a specific path. However, when using active measurement mechanisms in a multi-path topology, the default forwarding behavior is to go through one path. So, it cannot collect the information of all the paths at one time.

This document extends IOAM Trace Option with a multi-path flag to simplify multi-path IOAM active measurement, which promotes the information collection and topology restoration of a multi-path topology. It can help the operators to know the performance of network comprehensively and efficiently.

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 19 April 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	4
1.2. Terminology	4
2. IOAM extension	4
3. Multi-path Measurement Procedures	4
3.1. Encapsulating Node Procedures	5
3.2. Transit Node Procedures	5
3.2.1. Packet Operation	5
3.2.2. Packet Forwarding	5
3.3. Decapsulating Node Procedures	5
4. IANA Considerations	5
5. Security Considerations	6
6. References	6
6.1. Normative References	6
6.2. Informative References	7
Acknowledgements	7
Authors' Addresses	7

1. Introduction

In situ Operations, Administration, and Maintenance (IOAM) collects OAM information within the packet while the packet traverses a particular network domain. IOAM is used to complement mechanisms, such as Ping or Traceroute.

[RFC9322] defines the Active flag to indicate that a packet is used for active measurement and should be terminated by the decapsulating node. It also provides two active measurement use cases where the Active flag could be used.

However, active measurements are typically used to collect the information of a specific path, when using active measurement mechanisms in a multi-path topology (there are multiple paths from the source node to the destination node and ECMP, UCMP or other multi-path routing strategy is used.), the default forwarding behavior is to go through one path. So, it can't collect all the path's information from source node to destination node. An example of active measurement in a multi-path topology is shown as follow:

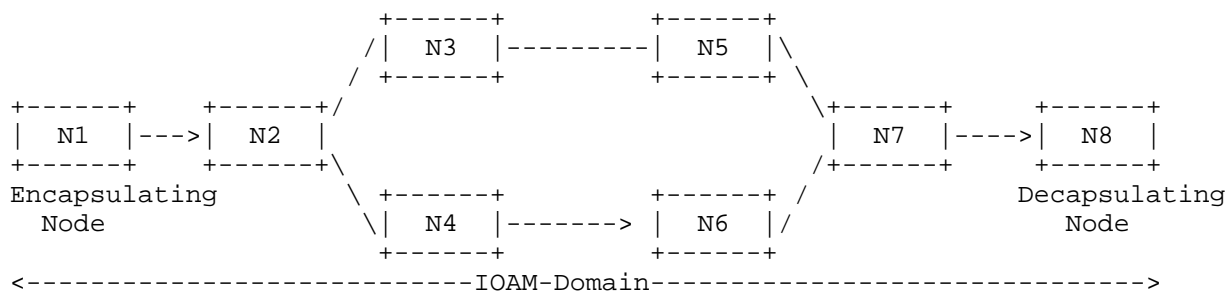


Figure 1: A multi-path topology

In Figure 1, node N1 is Encapsulating node, node N8 is the Decapsulating node, N2-N7 are transit node. Equal-Cost Multiple Path (ECMP) is applied in this topology. So, there are two paths form N1 to N8, one is N1-N2-N3-N5-N7-N8, and the other is N1-N2-N4-N6-N7-N8.

When N1 use IOAM probe packets or replicated data packets to measure the paths performance, the packets are forwarded along one of the paths (for example N1-N2-N4-N6-N7-N8), then the analyzer just can get one of the paths information, however the traffic packets are forwarded in all paths.

Although the IPv6 flow label and MPLS entropy label can be constructed variously according to the paths information to make packets go through all paths, but in some scenarios, it is hard to get all the available paths in advance.

This document extends IOAM Trace Option with a multi-path flag to simplify multi-path IOAM active measurement, which can promote the information collection and topology restoration of a multi-path topology without knowing all the available paths in advance.

1.1. 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.

1.2. Terminology

The abbreviations used in this document are:

ECMP: Equal-Cost Multiple Path

UCMP: Unequal-Cost Multiple Path

IOAM: In situ Operations, Administration, and Maintenance

2. IOAM extension

The format of IOAM Pre-allocated and Incremental Trace-Option header defined in [RFC9197] is shown as follows:




Figure 2: IOAM Pre-allocated and Incremental Trace-Option Header

This document defines a new flag X in the Flags field:

Bit X "Multipath" (M-bit): When set, the Multi-path flag indicates that when an active measurement packet arrives at a node which has multiple paths to the destination, the packet will be duplicated to every interface which can reach the destination.

3. Multi-path Measurement Procedures

This section describes the procedures of Encapsulating node, transit node, and Decapsulating node.

3.1. Encapsulating Node Procedures

The Encapsulating node should generate probe packets with a Trace Option that has its Active flag and Multipath flag set. The probe packets could be generated independently or replicated from some of the en-route data packets and should be terminated by the Decapsulating node.

3.2. Transit Node Procedures

3.2.1. Packet Operation

When the transit node receives a probe packet with the IOAM trace option, it should add its node ID, ingress interface ID, and egress interface ID to the IOAM trace option of the probe packet. For details about the content format, see section 4.4.2 of [RFC9197].

3.2.2. Packet Forwarding

When a transit node with multiple paths to the Destination node receives a packet with a Trace Option that has its Active flag and Multipath flag set, it SHOULD duplicate the packet to each egress interface that can reach the Destination node.

When the transit node has only one path to the destination node, it just needs to forward the packet it received to the egress interface.

If the Active flag is reset or the Multipath flag is reset, then the transit node MUST not duplicate the packet.

3.3. Decapsulating Node Procedures

When a Decapsulating node receives a probe packet with the IOAM trace option, it needs to add its node ID, ingress interface ID to the IOAM trace option of the probe packet. Then the Decapsulating node needs to export the IOAM data to the analyzer.

4. IANA Considerations

This document requests IANA to allocate a bit from the "IOAM Trace-Flags" registry:

Value	Description	Reference
Bit X	Multi-path Bit	This document

Table 1

5. Security Considerations

The security considerations of IOAM Active flag are discussed in [RFC9322], the solutions mitigating the attacks mentioned in [RFC9322] are also applicable in this document.

In addition, the duplication of probe packets may lead to other risks. When there is a loop in the topology, the probe packets may be replicated repeatedly. Even there is no loop in the topology, an attacker can replicate a lot of packets by setting the Multi-path flag in en-route packets, causing bandwidth degradation.

In order to mitigate the possible attacks, the IOAM enabled nodes should be able to:

- * Limit the generation rate of IOAM probe packets with the Multi-path flag.
- * Limit the maximum number of packet replication times, this could be realized by defining an Opaque State Snapshot filed in the Trace Option. The value of this field should decrease by one when a node duplicates the probe packet. The initial value of this field is set by the Encapsulating node, when its value decreases to zero, then the packet MUST not be duplicated anymore.

6. References

6.1. Normative References

- [RFC9322] Mizrahi, T., Brockners, F., Bhandari, S., Gafni, B., and M. Spiegel, "In Situ Operations, Administration, and Maintenance (IOAM) Loopback and Active Flags", RFC 9322, DOI 10.17487/RFC9322, November 2022, <<https://www.rfc-editor.org/rfc/rfc9322>>.
- [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/rfc/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/rfc/rfc8174>>.

6.2. Informative References

[RFC9197] Brockners, F., Ed., Bhandari, S., Ed., and T. Mizrahi, Ed., "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9197, DOI 10.17487/RFC9197, May 2022, <<https://www.rfc-editor.org/rfc/rfc9197>>.

Acknowledgements

The authors would like to thank Ernesto Ruffini, Thomas Graf, Greg Mirsky for the valuable comments to this work.

Authors' Addresses

Li Zhang
Huawei
China
Email: zhangli344@huawei.com

Tianran Zhou
Huawei
China
Email: zhoutianran@huawei.com

Junfeng Ma
CAICT
China
Email: majunfeng@caict.ac.cn