

ippm
Internet-Draft
Intended status: Standards Track
Expires: October 18, 2018

H. Song, Ed.
T. Zhou
Huawei
April 16, 2018

In-situ OAM Data Validation Option
draft-song-ippm-ioam-data-validation-option-02

Abstract

This document describes several potential performance scalability and capability issues when implementing in-situ OAM on heterogenous target network elements. The document proposes the corresponding solutions and modifications to the current in-situ OAM specification to mitigate the issues. Specifically, in-situ OAM is extended with data validation fields to cope with the node processing capability. We provide use cases to motivate our proposal and base the modifications on the current in-situ OAM header format specification.

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 October 18, 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	2
2. In-situ OAM Sampling and Data Validation	3
2.1. Valid Node Bitmap and Valid Data Bitmap	3
2.2. Use Cases	5
3. Security Considerations	6
4. IANA Considerations	6
5. Acknowledgments	6
6. Contributors	6
7. Informative References	6
Authors' Addresses	6

1. Introduction

In-situ OAM (iOAM) [I-D.brockners-inband-oam-requirements] records OAM information within user packets while the packets traverse a network. The data types and data formats for in-situ OAM data records have been defined in [I-D.ietf-ippm-ioam-data]. We identify several scalability issues for implementing the current iOAM specification and propose solutions in this draft.

iOAM can designate the flow to add the iOAM header and collect data on the flow forwarding path. The flow can have arbitrary granularity. However, processing the data can be a heavy burden for the network nodes, especially when some data needs to be calculated by the node (e.g., the transit delay). If the flow traffic is heavy, the node may not be able to handle the iOAM processing so many performance issues may occur, such as long latency and packet drop.

Although it is good for the OAM applications to gain the detailed information on every packet at every node, in many cases, such information is often repetitive and redundant. The large quantity of data would also burden the management plane which needs to collect and stream the data for analytics. It is also possible that some nodes cannot provide the requested data at all or are unwilling to provide some data for security or privacy concerns. So a trade-off is needed to balance the performance impact and the data availability and completeness.

We provide several motivating examples. To minimize the network impact, a network operator decides to collect the iOAM data only for initial and last flow packets (e.g., TCP packets with SYN, FIN, and RST flags).

In another example, a head node alternates two iOAM headers with each requesting a subset of iOAM data. Hence, each node on the flow path only needs to handle partial data. The requests can be balanced without exhausting the network nodes.

The above two cases can be realized by manipulating the iOAM header at the domain edge. It is also possible that a node is temporarily under heavy traffic load. It is in danger of dropping packets if it tries to satisfy all the iOAM data requests. It is also possible that, due to the privacy concern or capability issues, a node cannot satisfy the data request indicated in the iOAM header. In these cases, it would rather deny some requests than drop user traffic. This case can be realized by adding some auxiliary fields in the iOAM header.

More examples are listed in Section 2.2.

2. In-situ OAM Sampling and Data Validation

Based on the observation in Section 1, the source edge node should be able to define either the period or the probability to add the iOAM header to the selected flow packet. In this way, only a subset of the flow/sec packets would carry the OAM data, which not only reduces the overall iOAM data quantity but also reduces the processing work load of the network nodes.

Different data type bitmap templates can also be defined and used selectively. For example, template A includes a subset of data and template B includes another subset of data. The two templates can be used in the iOAM header for a flow alternately or in any predefined pattern. This is also an effective way to reduce the node processing load.

2.1. Valid Node Bitmap and Valid Data Bitmap

It is possible that even an iOAM capable node will not add data to the node data list as requested. In some cases, a node can be too busy to handle the data request or some types of the requested data is not available due to privacy and capability reasons. Therefore, we propose to add two bitmaps, a valid node bitmap and a valid data bitmap, to the iOAM specification.

The Node Valid Bitmap (NVB) is inserted before the Node Data List as shown in Figure 1. Each bit in the NVB corresponds to a hop on the packet's forwarding path. The bits are listed in the same order as the hop on the packet's forwarding path. The bitmap is set to all one at first. If a hop cannot add data to the Node Data List, the corresponding bit in the NVB is cleared to 0. The bit location for a

hop can be calculated from the length field (e.g, the bit index is equal to SSize-RHop).The valid node data items in the node data list is equal to the number of 1's in the NVB.

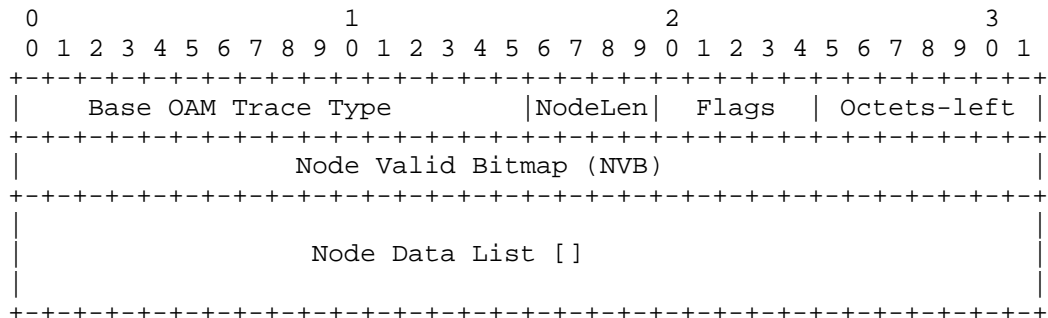


Figure 1: iOAM Header Format with Node Valid Bitmap (NVB)

NVB allows the head node to invalidate some nodes in advance. For example, if the head node wants to exclude the odd-numbered nodes from adding iOAM data, it can set all the corresponding bits to 0. Then at each node, if it finds its corresponding bit in the NVB is 0, it will simply skip the iOAM processing.

In addition to NVB, for each node data in the node data list, a Data Valid Bitmap (DVB) is added before the node data. The number of bits in the DVB is equal to the number of 1's in the OAM Trace Type bitmaps (excluding the next trace type bitmap indicator bits). When the bit is set, the corresponding data is valid in the node; otherwise, the corresponding data is invalid so the management plane should ignore it after the data is collected.

The size of the DVB can be padded to two or four bytes, which allow up to 16 or 32 types of data to be included in a node. The node data list format with the enhanced DVB is shown in Figure 2.

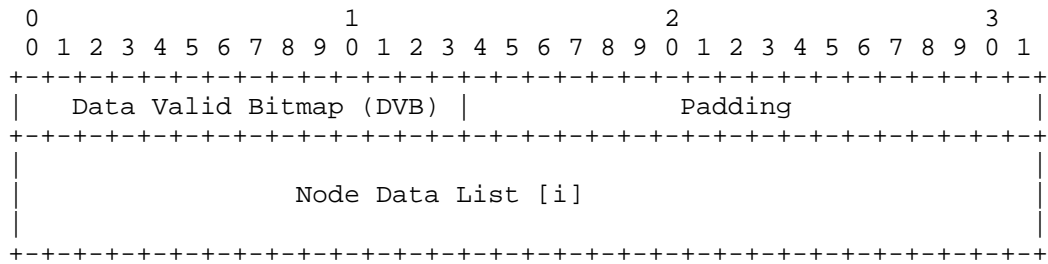


Figure 2: iOAM Node Data with Data Valid Bitmap (DVB)

2.2. Use Cases

We give some examples to show the usefulness of in-situ OAM sampling and data validation features.

- o An application needs to track a flow's forwarding path and knows the path will not change frequently, so it sets a low sampling rate to periodically insert the iOAM header to request the node ID.
- o In a heterogeneous data plane, some nodes support to provide data x but the other nodes do not support it. However, an application is still interested in collecting data x if available. In this case, iOAM header can still be configured to ask for data x but the nodes that cannot provide the data simply invalidates it by resetting the corresponding bit in the valid data bitmap.
- o Multiple sampling rate and multiple data request schema can be defined for a flow based on applications requirements and the data property, so for a flow packet, there can be no iOAM header or different iOAM headers. The node does not need to process all data all the time.
- o For security reason, a node decides to not participate in the iOAM data collection. While it processes the other iOAM header fields as usual, it does not set the node valid bit in the Node Valid Bitmap and add node data to the Node Data List.
- o To reduce the node processing load, the head node alternately uses two NVBs with one of them invalidating all the even-numbered nodes and the other invalidating all the odd-numbered nodes. Therefore, a node only needs to process the iOAM for every two packets of the flow.

3. Security Considerations

There is no extra security considerations beyond those have been identified by in-situ OAM protocol.

4. IANA Considerations

This memo includes no request to IANA.

5. Acknowledgments

We would like to thank Frank Brockners, Carlos Pignataro, Shwetha Bhandari, and Tal Mizrahi for helpful comments and suggestions.

6. Contributors

The document is inspired by numerous discussions with James N. Guichard. He also provided significant comments and suggestions to help improve this document.

7. Informative References

[I-D.brockners-inband-oam-requirements]

Brockners, F., Bhandari, S., Dara, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mozes, D., Mizrahi, T., <>, P., and r. remy@barefootnetworks.com, "Requirements for In-situ OAM", draft-brockners-inband-oam-requirements-02 (work in progress), October 2016.

[I-D.ietf-ippm-ioam-data]

Brockners, F., Bhandari, S., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., Chang, R., and d. daniel.bernier@bell.ca, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-00 (work in progress), September 2017.

Authors' Addresses

Haoyu Song (editor)
Huawei
2330 Central Expressway
Santa Clara, 95050
USA

Email: haoyu.song@huawei.com

Tianran Zhou
Huawei
156 Beiqing Road
Beijing, 100095
P.R. China

Email: zhoutianran@huawei.com