

ippm
Internet-Draft
Intended status: Standards Track
Expires: October 19, 2018

H. Song, Ed.
T. Zhou
Huawei
April 17, 2018

Control In-situ OAM Overhead with Segment IOAM
draft-song-ippm-segment-ioam-01

Abstract

This document describes a proposal which partitions an in-situ OAM (iOAM) domain into multiple segments in order to control the iOAM data overhead, adapt to the path MTU limitations, and enable new applications. We discuss several use cases to motivate our proposal and base the necessary 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 19, 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

Internet-Draft

Segment IOAM

April 2018

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Segment In-situ OAM	3
2.1.	Segment and Hops	3
2.2.	Considerations for Data Handling	4
2.3.	Use Cases	4
3.	Security Considerations	4
4.	IANA Considerations	5
5.	Acknowledgments	5
6.	Contributors	5
7.	Informative References	5
	Authors' Addresses	5

[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](#)].

iOAM may incur significant overhead on user packets. The overhead includes the iOAM header and the node data list for each network element.

The total size of data is limited by the MTU. When the number of required data types is large and the forwarding path length is long, it is possible that there is not enough space in the user packets to hold the iOAM header and data. The current proposal is to label the overflow status and stop adding new node data to the packet, leading to the loss of information.

Even if the header has enough space to hold the iOAM data, the overhead may be too large and consumes too much bandwidth. For example, if we assume moderate 20 bytes of data per node, a path with length of 10 will need 200 bytes to hold the data. This will inflate small 64-byte packets by more than four times. Even for the largest packet size (e.g., 1500 bytes), the overhead (>10%) is not negligible. Therefore, we need to limit the iOAM data overhead without sacrificing the data collection capability.

Here we have another interesting related issue. Packets can be dropped anywhere in a network for various reasons. If we can only collect iOAM data at the path end, we lose all data from the dropped

packets and have no idea where the packets are dropped. This defies the purpose of iOAM and makes those iOAM-enabled nodes work in vain.

2. Segment In-situ OAM

Based on the observation in [Section 1](#), we propose a method to limit the size of the node data list.

2.1. Segment and Hops

A hop is a node on a flow's forwarding path which is capable of processing iOAM data. A segment is a fixed number hops on a flow's forwarding path. While working in the "per hop" trace mode, the segment size (SSize) and the remaining hops (RHop), is added to the iOAM header at the edge. Initially, RHop is equal to SSize. At each hop, if RH is not zero, the node data is added to the node data list at the corresponding location and then RH is decremented by 1. If RH is equal to 0 when receiving the packet, the node needs to remove (in incremental trace option) or clear (in pre-allocated trace option) the iOAM node data list and reset RHop to SSize. Then the node will add its data to the node data list as if it is the edge node.

The stripped iOAM data at the segment edge can be immediately exported to a collector.

Figure 1 shows the proposed in-situ OAM header format. The bit 23 in the Flags field is used to indicate the current header is a segment iOAM header. In this context, the last octet in the iOAM header is partitioned into two 4-bit nibbles. The first nibble (SSize) is used to save the segment size and the second nibble (RHop) is used to save the remaining hops. This limits the maximum segment size to 15.

[illegible]




Figure 1: Segment iOAM Header Format

In the special case when SSize is set to 0, no data will be recorded in the node data list. The requested data listed in the OAM Trace

Type will be immediately exported to the collector. This way the iOAM overhead is minimized.

2.2. Considerations for Data Handling

At any hop when RHop is equal to 0, the node data list is copied from the iOAM header. The data can be encapsulated and reported to the controller or the edge node as configured. The encapsulation and report method is beyond the scope of this draft but should be comply with the method used by the iOAM edge node.

The actual size of the last segment may not be equal to SSize but this is not a problem.

2.3. Use Cases

Segment iOAM is necessary in the following example scenarios:

- o Segment iOAM can be used to detect at which segment the flow packet is dropped. If the SSize is set to 1, then the exact drop node can be identified. The iOAM data before the dropping point is also retained.
- o The path MTU allows to add at most k node data in the list to avoid fragmentation. Therefore SSize is set to k and at each hop where RHop is 0, the node data list is retrieved and sent in a standalone packet.
- o A flow contains mainly short packets and travels a long path. It

would be inefficient to keep a large node data list in the packet so the network bandwidth utilization rate is low. In this case, segment iOAM can be used to limit the ratio of the iOAM data to the flow packet payload.

- o The network allows at most n bytes budget for the iOAM data. There is a tradeoff between the number of data types that can be collected and the number of hops for data collecting. The segment size is therefore necessary to meet the application's data requirement (i.e., $SSize * Node\ Data\ Size < n$).

[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, and Shwetha Bhandari 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

Song & Zhou

Expires October 19, 2018

[Page 5]

Internet-Draft

Segment IOAM

April 2018

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

Email: zhoutianran@huawei.com

