ippm
Internet-Draft
Intended status: Standards Track
Expires: May 6, 2021

B. Gafni
Nvidia
H. Liu
R. Miao
Alibaba Group
M. Spiegel
Barefoot Networks, an Intel company

November 02, 2020

Additional data fields for IOAM Trace Option Types
draft-gafni-ippm-ioam-additional-data-fields-00

Abstract

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a path between two points in the network. This document discusses additional data fields and associated data types to be added to the IOAM data fileds described in [I-D.ietf-ippm-ioam-data]. In-situ OAM data fields can be encapsulated into a variety of protocols such as NSH, Segment Routing, Geneve, IPv6 (via extension header), or IPv4.

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 May 6, 2021.

Copyright Notice

Table of Contents

1. Introduction

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a path between two points in the network. This document is adding additional data fields that can be reported by the network as part of IOAM.

2. Conventions

2.1. 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.

2.2. Abbreviations

Abbreviations used in this document:

IOAM: In-situ Operations, Administration, and Maintenance

3. Additional Data Fields

This draft extends [I-D.ietf-ippm-ioam-data] with additional data fields. The additional suggested data fields are:

- Transmitted Bytes from an interface
- Speed of an interface
- Interface errors

The addition of these new data fields is intended to help network operators to better manage their networks, where more data is requried with regards to the activity and quality of the network ports. For example, one framework that may take advantage of these new data fileds is HPCC, which is proposed at [I-D.pan-tsvwg-hpccplus]. This section discusses the needed ammendments to the IOAM Trace header and the format of the added data fields themselves.

3.1. IOAM Trace Option-Types Ammendments

IOAM Trace Option-Types and their headers are defined in section 4.4 of [I-D.ietf-ippm-ioam-data]. As shown in section 4.4.1, the trace option header includes an IOAM-Trace-Type which is a "A 24-bit identifier which specifies which data types are used in this node data list". In order to extend [I-D.ietf-ippm-ioam-data] it is required to allocate respective bits specifying the additional data fields to be added to the packet. This draft is asking for the allocation of additional 2 bits:

Bit 12 When set indicates presence of Transmitted Bytes from an interface.

Bit 13 When set indicates presence of Speed of an interface and Interface errors.

Section 3.2 describes the new suggested data types and their formats.

3.2. The Additional IOAM Node Data Fields and Associated Formats

The Data fields and associated data types for each of the additional IOAM Data Fields are shown below:

Transmitted Bytes from an interface: 4-octet field defined as follows:

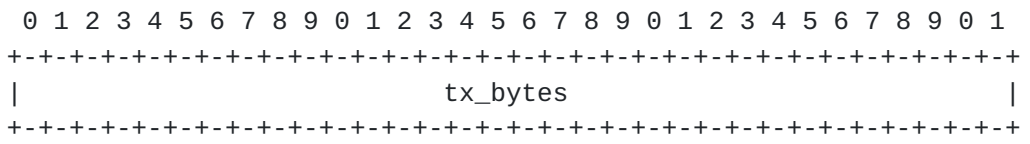

tx_bytes: 4-octet unsigned integer field. This field indicates how many bytes have been transmitted from the egress interface the packet is going out from. Note that this field may wrap around. As an example, for a 100Gbps port this field may wrap around within less than 3 seconds. This field is usable to determine the amount of data going through the path a flow is going through. Following multiple packets traversing the same interface, together with a timestamp, allows a network operator to gauge the amount of traffic going through the interface in total and relative to the flow it tracks. This data in turn may help to better control the traffic and take decisions related to the performance of the flow and the network.

Speed of an interface and Total errors of an interface: 4-octet field defined as follows:

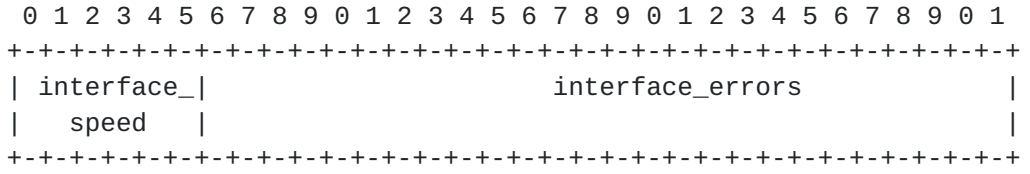

interface_speed: 6 bits unsigned integer field. This field indicates the current operational speed of the interface. The procedure to allocate, manage and map the interface_speed values into the actual speed is beyond the scope of this document. This field is usable to detect whether a packet or a flow is going through a path which has enough capacity compared to the expectation of the operator. Changes in the speed of the connectivty may require changing routing decisions or troubleshooting the links under consideration. When an operator intends to take a decision about the amount of data to transmit per flow, this data is helpful as well to track.

interface_errors: 26 bits unsigned integer field. This field inciates how many errors, such as packet drops due to CRC errors, have been detected on the interface used to deliver the packet. This data is helpful in order to understand the risk associated with the packet, or the flow it belongs to, as it shows the quality of the interfaces it uses as part of its path in the network. It can also point out potential issues that other packets from the same flow might have experienced.

4. Security Considerations

TBD

5. IANA Considerations

TBD

6. References

6.1. 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>.

6.2. Informative References

[I-D.ietf-ippm-ioam-data] Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", draft-ietf-ippm-ioam-data-10 (work in progress), July 2020.

[I-D.pan-tsvwg-hpccplus] Miao, R., Liu, H., Pan, R., Lee, J., Kim, C., Gafni, B., and Y. Shpigelman, "HPCC++: Enhanced High Precision Congestion Control", draft-pan-tsvwg-hpccplus-02 (work in progress), September 2020.

Authors' Addresses

Barak Gafni
Nvidia
350 Oakmead Parkway, Suite 100
Sunnyvale, CA 94085
U.S.A.

Email: gbarak@nvidia.com

Hongqiang H. Liu
Alibaba Group
108th Ave NE, Suite 800
Bellevue, WA 98004
U.S.A.

Email: hongqiang.liu@alibaba-inc.com

Rui Miao
Alibaba Group
525 Almanor Ave, 4th Floor
Sunnyvale, CA 94085
USA

Email: miao.rui@alibaba-inc.com

Mickey Spiegel
Barefoot Networks, an Intel
      company
4750 Patrick Henry Drive
Santa Clara, CA 95054
US

Email: mickey.spiegel@intel.com