

Workgroup: Network Working Group
Internet-Draft:
draft-tgraf-opsawg-ipfix-inband-telemetry-01
Published: 28 July 2022
Intended Status: Standards Track
Expires: 29 January 2023
Authors: T. Graf B. Claise A. Huang Feng
 Swisscom Huawei INSA-Lyon
Export of Forwarding Path Delay in IPFIX

Abstract

This document introduces new IP Flow Information Export (IPFIX) information elements to expose the Inband Telemetry measured forwarding path delay on the IOAM transit and decapsulation nodes.

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 29 January 2023.

Copyright Notice

Copyright (c) 2022 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. | Performance Metrics |
| 2.1. | IP One-Way Delay Hybrid Type I Passive Registry Entries |
| 2.1.1. | Summary |
| 2.1.2. | Description |
| 2.1.3. | Change Controller |
| 2.1.4. | Version of Registry Format |
| 2.2. | Metric Definition |
| 2.2.1. | Reference Definition |
| 2.2.2. | Fixed Parameters |
| 2.3. | Method of Measurement |
| 2.3.1. | Reference Methods |
| 2.3.2. | Packet Stream Generation |
| 2.3.3. | Traffic Filtering (Observation) Details |
| 2.3.4. | Sampling Distribution |
| 2.3.5. | Runtime Parameters and Data Format |
| 2.3.6. | Roles |
| 2.4. | Output |
| 2.4.1. | Type |
| 2.4.2. | Reference Definition |
| 2.4.3. | Administrative Items |
| 2.4.4. | Comments and Remarks |
| 3. | IPFIX Information Elements |
| 4. | Use Cases |
| 5. | IANA Considerations |
| 5.1. | Performance Metrics |
| 5.2. | IPFIX Entities |
| 5.2.1. | PathDelayMeanDeltaMicroseconds |
| 5.2.2. | PathDelayMeanDeltaNanoseconds |
| 5.2.3. | PathDelayMinDeltaMicroseconds |
| 5.2.4. | PathDelayMinDeltaNanoseconds |
| 5.2.5. | PathDelayMaxDeltaMicroseconds |
| 5.2.6. | PathDelayMaxDeltaNanoseconds |
| 5.2.7. | PathDelaySumDeltaMicroseconds |
| 5.2.8. | PathDelaySumDeltaNanoseconds |
| 6. | Operational Considerations |
| 6.1. | Time Accuracy |
| 6.2. | Mean Delay |
| 6.3. | IOAM Packet Time Stamps |
| 7. | Security Considerations |
| 8. | Acknowledgements |
| 9. | References |
| 9.1. | Normative References |
| 9.2. | Informative References |
| | Authors' Addresses |

1. Introduction

Network operators want a statistical delay view of their networks. They want to understand where in the network, for which customer traffic, how much and why delay is being accumulated. In order to answer why and where, delay needs to be reported into device and control-plane context. In order to understand which customer traffic is affected, delay needs to be reported into customer data-plane context. That enables network operators to quickly identify when the control-plane updates the current path with a different next-hop and therefore the forwarding path changes to different nodes and interfaces, how the path delay changes for which customer traffic.

With Inband Telemetry, defined in [In-situ OAM](#) [[I-D.ietf-ippm-ioam-deployment](#)], [Path Tracing](#) [[I-D.filsfils-spring-path-tracing](#)] and [In-situ Flow Information Telemetry](#) [[I-D.song-opsawg-ifit-framework](#)], the path delay between two endpoints is measured by inserting a timestamp in the packet.

Inband Telemetry can be distinguished between two modes. Passport mode, [[RFC9197](#)], where only the last hop in the forwarding path of the Inband Telemetry domain exposes all the metrics, and postcard mode, [[I-D.song-ippm-postcard-based-telemetry](#)], where the metrics are also exposed in the transit nodes. In both modes the forwarding path exposes performance metrics allowing to determine how much delay has been accumulated on which hop.

This document defines eight new IPFIX Information Elements (IEs), exposing the forwarding path delay on IOAM transit and decapsulation nodes. Since these IPFIX IEs are performance metrics [[RFC8911](#)], they must be registered as registered performance metrics [[RFC8911](#)] in the ["IANA Performance Metric Registry](#) [[IANA-PERF-METRIC](#)].

| Performance Metric | IPFIX Entity |
|-------------------------------|---------------------------------------|
| OWDelay_HybridType1_Passive_I | PathDelayMeanDeltaMicroseconds (TBD5) |
| P_RFCTBD_Seconds_Mean (TBD1) | PathDelayMeanDeltaNanoseconds (TBD6) |
| OWDelay_HybridType1_Passive_I | PathDelayMinDeltaMicroseconds (TBD7) |
| P_RFCTBD_Seconds_Min (TBD2) | PathDelayMinDeltaNanoseconds (TBD8) |
| OWDelay_HybridType1_Passive_I | PathDelayMaxDeltaMicroseconds (TBD9) |
| P_RFCTBD_Seconds_Max (TBD3) | PathDelayMaxDeltaNanoseconds (TBD10) |
| OWDelay_HybridType1_Passive_I | PathDelaySumDeltaMicroseconds (TBD11) |
| P_RFCTBD_Seconds_Sum (TBD4) | PathDelaySumDeltaNanoseconds (TBD12) |

Table 1: Correspondance between IE and performance metric registry

The delay is measured by calculating the difference between the timestamp imposed with Inband Telemetry in the packet at the IOAM encapsulation node and the timestamp exported in the IPFIX flow record from the IOAM transit and decapsulation nodes. Depending on the IE, the lowest, highest or the sum of measured path delay is being exported.

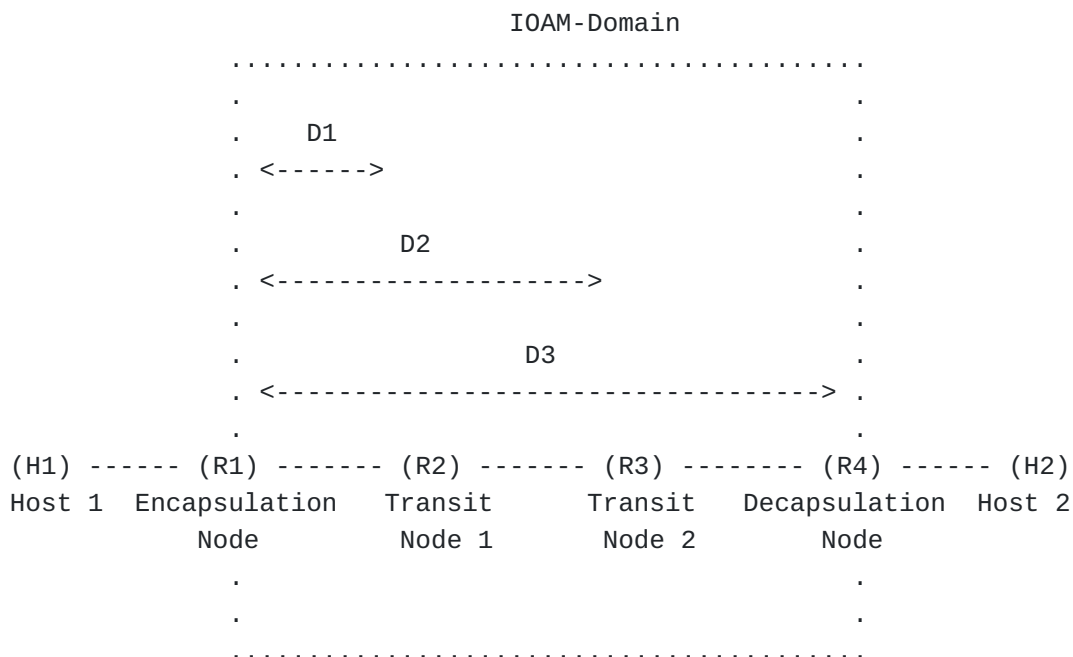


Figure 1: IOAM Delay use case. Packets flow from host 1 to host 2.

On the usecase showed in [Figure 1](#) using IOAM to export the delay metrics, the node R2 exports the delay D1, the node R3 exports the delay D2 and the decapsulation node R4 exports the total delay D3 using IPFIX.

2. Performance Metrics

This section defines and describes the new performance metrics

2.1. IP One-Way Delay Hybrid Type I Passive Registry Entries

This section specifies four Registry Entries for the Hybrid Type I Passive assessment of IP One-Way Delay.

All column entries besides the ID, Name, Description, and Output Reference Method categories are the same; thus, this section defines four closely related Registry Entries. As a result, IANA has assigned corresponding URLs to each of the four Named Metrics.

2.1.1. Summary

This category includes multiple indexes to the Registry Entry: the element ID and Metric Name.

2.1.1.1. ID (Identifier)

<insert a numeric Identifier, an integer, TBD>

2.1.1.2. Name

IANA has allocated the numeric Identifiers TBD1-4 for the four Named Metric Entries in this section

2.1.1.3. Name

TBD1: OWDelay_HybridType1_Passive_IP_RFCTBD_Seconds_Mean

TBD2: OWDelay_HybridType1_Passive_IP_RFCTBD_Seconds_Min

TBD3: OWDelay_HybridType1_Passive_IP_RFCTBD_Seconds_Max

TBD4: OWDelay_HybridType1_Passive_IP_RFCTBD_Seconds_Sum

2.1.1.4. URI

URL: https://www.iana.org/assignments/performance-metrics/OWDelay_HybridType1_Passive_IP_RFCTBD_Seconds_Mean

URL: https://www.iana.org/assignments/performance-metrics/OWDelay_HybridType1_Passive_IP_RFCTBD_Seconds_Min

URL: https://www.iana.org/assignments/performance-metrics/OWDelay_HybridType1_Passive_IP_RFCTBD_Seconds_Max

URL: https://www.iana.org/assignments/performance-metrics/OWDelay_HybridType1_Passive_IP_RFCTBD_Seconds_Sum

2.1.2. Description

This metric assesses the one-way delay of IP packets constituting a single connection between two hosts. We consider the measurement of one-way delay based on a single Observation Point (OP) [RFC7011] somewhere in the network. The output is the one-way delay for all successfully forwarded packets expressed as the <statistic> of their conditional delay distribution, where <statistic> is one of:

*Mean

*Min

*Max

*Sum

2.1.3. Change Controller

IETF

2.1.4. Version of Registry Format

1.0

2.2. Metric Definition

This category includes columns to prompt the entry of all necessary details related to the metric definition, including the immutable document reference and values of input factors, called "Fixed Parameters".

2.2.1. Reference Definition

Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, RFC 7679, DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>. [RFC7679]

Morton, A. and E. Stephan, "Spatial Composition of Metrics", RFC 6049, DOI 10.17487/RFC6049, January 2011, <<https://www.rfc-editor.org/info/rfc6049>>. [RFC6049]

Section 3.4 of [[RFC7679](#)] provides the reference definition of the singleton (single value) one-way delay metric. Section 4.4 of [[RFC7679](#)] provides the reference definition expanded to cover a multi-value sample. Note that terms such as "singleton" and "sample" are defined in section 2 of [[RFC2330](#)].

With the OP [[RFC7011](#)] typically located between the hosts participating in the IP connection, the one-way delay metric requires one individual measurement between the OP and sourcing host, such that the Spatial Composition [[RFC6049](#)] of the measurements yields a one-way delay singleton.

2.2.2. Fixed Parameters

Traffic Filters:

IPv4 header values:

DSCP: Set to 0

IPv6 header values:

DSCP: Set to 0

Hop Count: Set to 255

Flow Label: Set to 0

Extension Headers: None

2.3. Method of Measurement

This category includes columns for references to relevant sections of the RFC(s) and any supplemental information needed to ensure an unambiguous method for implementations.

2.3.1. Reference Methods

The foundational methodology for this metric is defined in section 4 of [[RFC7323](#)] using the Timestamps option with modifications that allow application at a mid-path OP [[RFC7011](#)].

The Traffic Filter at the OP is configured to observe a single IP connection.

2.3.2. Packet Stream Generation

N/A

2.3.3. Traffic Filtering (Observation) Details

The Fixed Parameters above give a portion of the Traffic Filter. Other aspects will be supplied as Runtime Parameters (below).

2.3.4. Sampling Distribution

This metric requires a partial sample of all packets that qualify according to the Traffic Filter criteria.

2.3.5. Runtime Parameters and Data Format

Runtime Parameters are input factors that must be determined, configured into the measurement system, and reported with the results for the context to be complete.

Src: The IP address of the host in the host A Role (format ipv4-address-no-zone value for IPv4 or ipv6-address-no-zone value for IPv6; see section 4 of [\[RFC6991\]](#)).

Dst: The IP address of the host in the host B Role (format ipv4-address-no-zone value for IPv4 or ipv6-address-no-zone value for IPv6; see section 4 of [\[RFC6991\]](#)).

TTL or Hop Limit: Set at desired value.

DSCP: Set at desired value.

IPv6 Flow Label: Set at desired value.

Timestamp: The timestamp when the packet is being received at IOAM encapsulation node. Format depends on Inband Telemetry implementation. For IOAM, Section 4.4.1 of [\[RFC9197\]](#) describes what kind of timestamps are supported. Section 4.4.2.3 and 4.4.2.4 describe where the timestamp is being inserted. For Path Tracing, Section 4.1 of [\[I-D.filsfils-spring-path-tracing\]](#) describes what kind of timestamps are supported. Section 9.2 describe the SRH path tracing TLV where the timestamp is being inserted.

2.3.6. Roles

host A: Launches the IP packet to open the connection. The Role of "host A" is synonymous with the IP address used at host A.

host B: Receives the IP packet to open the connection. The Role of "host B" is synonymous with the IP address used at host B.

Encapsulation Node: Receives the IP packet to open the connection and encapsulates the timestamp into the packet. The Role of

"Encapsulation Node" is synonymous with the timestamp inserted in the packet.

Transit Node: Receives the IP packet to open the connection and measures the delay between the timestamp in the packet and the timestamp when the packet was received.

Decapsulation Node: Receives the IP packet to open the connection and measures the delay between the timestamp in the packet and the timestamp when the packet was received and removes the IOAM header from the packet.

2.4. Output

This category specifies all details of the output of measurements using the metric.

2.4.1. Type

OWDelay Types are discussed in the subsections below.

2.4.2. Reference Definition

For all output types:

OWDelay_HybridType1_Passive_IP: The one-trip delay of one IP packet is a Singleton

For each <statistic>, Singleton one of the following subsections applies.

2.4.2.1. Mean

The mean SHALL be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded) -- a single value, as follows:

See section 4.1 of [[RFC3393](#)] for details on the conditional distribution to exclude undefined values of delay, and see section 5 of [[RFC6703](#)] for background on this analysis choice.

See section 4.2.2 of [[RFC6049](#)] for details on calculating this statistic; see also section 4.2.3 of [[RFC6049](#)].

Mean: The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see section 9.3 of [[RFC6020](#)]) with a resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per section 6 of [[RFC5905](#)].

2.4.2.2. Min

The minimum SHALL be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded) -- a single value, as follows:

See section 4.1 of [[RFC3393](#)] for details on the conditional distribution to exclude undefined values of delay, and see section 5 of [[RFC6703](#)] for background on this analysis choice.

See section 4.3.2 of [[RFC6049](#)] for details on calculating this statistic; see also section 4.3.3 of [[RFC6049](#)].

Min: The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see section 9.3 of [[RFC6020](#)]) with a resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per section 6 of [[RFC5905](#)].

2.4.2.3. Max

The maximum SHALL be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded) -- a single value, as follows:

See section 4.1 of [[RFC3393](#)] for details on the conditional distribution to exclude undefined values of delay, and see section 5 of [[RFC6703](#)] for background on this analysis choice.

See section 4.3.2 of [[RFC6049](#)] for a closely related method for calculating this statistic; see also section 4.3.3 of [[RFC6049](#)]. The formula is as follows:

Max = (FiniteDelay[j])
such that for some index, j, where $1 \leq j \leq N$
FiniteDelay[j] >= FiniteDelay[n] for all n

Max: The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see section 9.3 of [[RFC6020](#)]) with a resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per section 6 of [[RFC5905](#)].

2.4.2.4. Sum

The sum SHALL be calculated using the conditional distribution of all packets with a finite value of one-way delay (undefined delays are excluded) -- a single value, as follows:

See section 4.1 of [[RFC3393](#)] for details on the conditional distribution to exclude undefined values of delay, and see section 5 of [[RFC6703](#)] for background on this analysis choice.

See section 4.3.5 of [[RFC6049](#)] for details on calculating this statistic. However in this case FiniteDelay or MaxDelay MAY be used.

Sum: The time value of the result is expressed in units of seconds, as a positive value of type decimal64 with fraction digits = 9 (see section 9.3 of [[RFC6020](#)]) with a resolution of 0.000000001 seconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per section 6 of [[RFC5905](#)].

2.4.2.5. Metric Units

The <statistic> of one-way delay is expressed in seconds, where <statistic> is one of:

*Mean

*Min

*Max

*Sum

The one-way delay of the IP connection singleton is expressed in seconds.

2.4.2.6. Calibration

Passive Measurements at an OP could be calibrated against an Active Measurement at host A where the Active Measurement represents the ground truth.

2.4.3. Administrative Items

2.4.3.1. Status

Current

2.4.3.2. Requester

This RFC

2.4.3.3. Revision

1.0

2.4.3.4. Revision Date

RFC Date

2.4.4. Comments and Remarks

None

3. IPFIX Information Elements

This section defines and describes the new IPFIX IEs.

PathDelayMeanDeltaMicroseconds

16-bit unsigned integer that identifies the mean path delay in microseconds, between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node).

PathDelayMeanDeltaNanoseconds

32-bit unsigned integer that identifies the mean path delay in nanoseconds, between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node).

PathDelayMinDeltaMicroseconds

16-bit unsigned integer that identifies the lowest path delay in microseconds, between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node).

PathDelayMinDeltaNanoseconds

32-bit unsigned integer that identifies the lowest path delay in nanoseconds, between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node).

PathDelayMaxDeltaMicroseconds

16-bit unsigned integer that identifies the highest path delay in microseconds, between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node).

PathDelayMaxDeltaNanoseconds

32-bit unsigned integer that identifies the highest path delay in nanoseconds, between the IOAM encapsulation node and the local

node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node).

PathDelaySumDeltaMicroseconds

32-bit unsigned integer that identifies the sum of the path delay in microseconds, between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node).

PathDelaySumDeltaNanoseconds

64-bit unsigned integer that identifies the sum of the path delay in nanoseconds, between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node).

4. Use Cases

The measured forwarding path delay can be aggregated with Flow Aggregation as defined in [[RFC7015](#)] to the following device and control-plane dimensions to determine:

- *With node id and egressInterface(IE14), on which node which logical egress interfaces have been contributing to how much delay.
- *With node id and egressPhysicalInterface(253), on which node which physical egress interfaces have been contributing to how much delay.
- *With ipNextHopIPv4Address(IE15) or ipNextHopIPv6Address(IE62), the forwarding path to which next-hop IP contributed to how much delay.
- *With mplsTopLabelIPv4Address(IE47) or srhActiveSegmentIPv6 from [[I-D.tgraf-opsawg-ipfix-srv6-srh](#)], the forwarding path to which MPLS top label IPv4 address or SRv6 active segment contributed to how much delay.
- *BGP communities are often used for setting a path priority or service selection. With bgpDestinationExtendedCommunityList(488) or bgpDestinationCommunityList(485) or bgpDestinationLargeCommunityList(491) which group of prefixes accumulated at which node how much delay.
- *With destinationIPv4Address(13), destinationTransportPort(11), protocolIdentifier (4) and sourceIPv4Address(IE8), the forwarding path delay on each node from each IPv4 source address to a specific application in the network.

Taking figure 1 from section 1 as topology example. Below example table shows the aggregated delay per each node, egressInterface and srhActiveSegmentIPv6.

| Path Delay | Node | egressInterface | srhActiveSegmentIPv6 |
|------------|------|-----------------|----------------------|
| 0 ns | R1 | 276 | 2001:db8::4 |
| 3122 ns | R2 | 312 | 2001:db8::4 |
| 4432 ns | R3 | 27 | 2001:db8::4 |
| 7237 ns | R4 | 854 | 2001:db8::4 |

Table 2: Example table of measured delay. Ascending by delay.

5. IANA Considerations

5.1. Performance Metrics

This document requests IANA to create new performance metrics under the "Performance Metrics" registry [[RFC8911](#)] with the values defined in section 2.

5.2. IPFIX Entities

This document requests IANA to create new IEs (see table 3) under the "IPFIX Information Elements" registry [[RFC7012](#)] available at ["IANA Performance Metric Registry" \[IANA-PERF-METRIC\]](#) and assign the following initial code points.

| Element | Name |
|---------|--------------------------------|
| ID | |
| TBD5 | PathDelayMeanDeltaMicroseconds |
| | |
| TBD6 | PathDelayMeanDeltaNanoseconds |
| | |
| TBD7 | PathDelayMinDeltaMicroseconds |
| | |
| TBD8 | PathDelayMinDeltaNanoseconds |
| | |
| TBD9 | PathDelayMaxDeltaMicroseconds |
| | |
| TBD10 | PathDelayMaxDeltaNanoseconds |
| | |
| TBD11 | PathDelaySumDeltaMicroseconds |
| | |
| TBD12 | PathDelaySumDeltaNanoseconds |
| | |

Table 3: Creates IEs in the "IPFIX Information Elements" registry

Note to the RFC-Editor:

*Please replace TBD5 - TBD12 with the values allocated by IANA

*Please replace the [RFC-to-be] with the RFC number assigned to this document

5.2.1. PathDelayMeanDeltaMicroseconds

Name: PathDelayMeanDeltaMicroseconds ElementID: TBD5 Description:
This Information Element identifies the mean path delay between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node) in microseconds. Abstract Data Type: unsigned16 Data Type Semantics: OctedDelta Reference: [RFC-to-be], xxx

5.2.2. PathDelayMeanDeltaNanoseconds

Name: PathDelayMeanDeltaNanoseconds ElementID: TBD6 Description: This Information Element identifies the mean path delay between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node) in nanoseconds. Abstract Data Type: unsigned32 Data Type Semantics: OctedDelta Reference: [RFC-to-be], xxx

5.2.3. PathDelayMinDeltaMicroseconds

Name: PathDelayMinDeltaMicroseconds ElementID: TBD7 Description: This Information Element identifies the lowest path delay between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node) in microseconds. Abstract Data Type: unsigned16 Data Type Semantics: OctedDelta Reference: [RFC-to-be], xxx

5.2.4. PathDelayMinDeltaNanoseconds

Name: PathDelayMinDeltaNanoseconds ElementID: TBD8 Description: This Information Element identifies the lowest path delay between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node) in nanoseconds. Abstract Data Type: unsigned32 Data Type Semantics: OctedDelta Reference: [RFC-to-be], xxx

5.2.5. PathDelayMaxDeltaMicroseconds

Name: PathDelayMaxDeltaMicroseconds ElementID: TBD9 Description: This Information Element identifies the highest path delay between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node) in microseconds. Abstract Data Type: unsigned16 Data Type Semantics: OctedDelta Reference: [RFC-to-be], xxx

5.2.6. PathDelayMaxDeltaNanoseconds

Name: PathDelayMaxDeltaNanoseconds ElementID: TBD10 Description: This Information Element identifies the highest path delay between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node) in nanoseconds. Abstract Data Type: unsigned32 Data Type Semantics: OctedDelta Reference: [RFC-to-be], xxx

5.2.7. PathDelaySumDeltaMicroseconds

Name: PathDelaySumDeltaMicroseconds ElementID: TBD11 Description: This Information Element identifies the sum of the path delay between the IOAM encapsulation node and the local node with the IOAM

domain (either an IOAM transit node or an IOAM decapsulation node) in microseconds. Abstract Data Type: unsigned32 Data Type Semantics: OctedDelta Reference: [RFC-to-be], xxx

5.2.8. PathDelaySumDeltaNanoseconds

Name: PathDelaySumDeltaNanoseconds ElementID: TBD12 Description: This Information Element identifies the sum of the path delay between the IOAM encapsulation node and the local node with the IOAM domain (either an IOAM transit node or an IOAM decapsulation node) in nanoseconds. Abstract Data Type: unsigned64 Data Type Semantics: OctedDelta Reference: [RFC-to-be], xxx

6. Operational Considerations

6.1. Time Accuracy

The same recommendation as defined in section 4.5 of [[RFC5153](#)] for IPFIX applies in terms of clock precision to this document as well.

6.2. Mean Delay

The mean (average) path delay can be calculated by dividing the PathDelaySumDeltaMicroseconds(TBD5) or PathDelaySumDeltaNanoseconds(TBD6) by the packetDeltaCount(2) at the IPFIX data collection.

6.3. IOAM Packet Time Stamps

For IOAM, Section 4.4.1 of [[RFC9197](#)] describes what kind of timestamps are supported. Section 4.4.2.3 and 4.4.2.4 describe where the timestamp is being inserted.

For Path Tracing, Section 4.1 of [[I-D.filsfils-spring-path-tracing](#)] describes what kind of timestamps are supported. Section 9.2 describe the SRH path tracing TLV where the timestamp is being inserted.

7. Security Considerations

There are no significant extra security considerations regarding the allocation of these new IPFIX IEs compared to [[RFC7012](#)].

8. Acknowledgements

The authors would like to thank xxx for their review and valuable comments.

9. References

9.1. Normative References

[IANA-PERF-METRIC] "IANA Performance Metric Registry", <<https://www.iana.org/assignments/performance-metrics/performance-metrics.xhtml>>.

[RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, DOI 10.17487/RFC7012, September 2013, <<https://www.rfc-editor.org/info/rfc7012>>.

[RFC8911] Bagnulo, M., Claise, B., Eardley, P., Morton, A., and A. Akhter, "Registry for Performance Metrics", DOI 10.17487/RFC8911, RFC 8911, November 2021, <<https://www.rfc-editor.org/info/rfc8911>>.

9.2. Informative References

[I-D.filsfils-spring-path-tracing]

Filsfils, C., Abdelsalam, A., Garvia, P. C., Yufit, M., Graf, T., Su, Y., Matsushima, S., and M. Valentine, "Path Tracing in SRv6 networks", Work in Progress, Internet-Draft, draft-filsfils-spring-path-tracing-01, 30 May 2022, <<https://www.ietf.org/archive/id/draft-filsfils-spring-path-tracing-01.txt>>.

[I-D.ietf-ippm-ioam-deployment] Brockners, F., Bhandari, S., Bernier, D., and T. Mizrahi, "In-situ OAM Deployment", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-deployment-01, 11 April 2022, <<https://www.ietf.org/archive/id/draft-ietf-ippm-ioam-deployment-01.txt>>.

[I-D.song-ippm-postcard-based-telemetry]

Song, H., Mirsky, G., Filsfils, C., Abdelsalam, A., Zhou, T., Li, Z., Mishra, G., Shin, J., and K. Lee, "In-Situ OAM Marking-based Direct Export", Work in Progress, Internet-Draft, draft-song-ippm-postcard-based-telemetry-12, 12 May 2022, <<https://www.ietf.org/archive/id/draft-song-ippm-postcard-based-telemetry-12.txt>>.

[I-D.song-opsawg-ifit-framework] Song, H., Qin, F., Chen, H., Jin, J., and J. Shin, "A Framework for In-situ Flow Information Telemetry", Work in Progress, Internet-Draft, draft-song-opsawg-ifit-framework-17, 22 February 2022,

<<https://www.ietf.org/archive/id/draft-song-opsawg-ifit-framework-17.txt>>.

- [I-D.tgraf-opsawg-ipfix-srv6-srh] Graf, T., Claise, B., and P. Francois, "Export of Segment Routing IPv6 Information in IP Flow Information Export (IPFIX)", Work in Progress, Internet-Draft, draft-tgraf-opsawg-ipfix-srv6-srh-05, 24 July 2022, <<https://www.ietf.org/archive/id/draft-tgraf-opsawg-ipfix-srv6-srh-05.txt>>.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, DOI 10.17487/RFC2330, May 1998, <<https://www.rfc-editor.org/info/rfc2330>>.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, DOI 10.17487/RFC3393, November 2002, <<https://www.rfc-editor.org/info/rfc3393>>.
- [RFC5153] Boschi, E., Mark, L., Quittek, J., Stiemerling, M., and P. Aitken, "IP Flow Information Export (IPFIX) Implementation Guidelines", RFC 5153, DOI 10.17487/RFC5153, April 2008, <<https://www.rfc-editor.org/info/rfc5153>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.
- [RFC6049] Morton, A. and E. Stephan, "Spatial Composition of Metrics", RFC 6049, DOI 10.17487/RFC6049, January 2011, <<https://www.rfc-editor.org/info/rfc6049>>.
- [RFC6703] Morton, A., Ramachandran, G., and G. Maguluri, "Reporting IP Network Performance Metrics: Different Points of

View", DOI 10.17487/RFC6703, RFC 6703, August 2012,
<<https://www.rfc-editor.org/info/rfc6703>>.

[RFC6991] Schoenwaelder, J., Ed., "Common YANG Data Types", RFC 6991, DOI 10.17487/RFC6991, July 2013, <<https://www.rfc-editor.org/info/rfc6991>>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.

[RFC7015] Trammell, B., Wagner, A., and B. Claise, "Flow Aggregation for the IP Flow Information Export (IPFIX) Protocol", DOI 10.17487/RFC7015, RFC 7015, September 2013, <<https://www.rfc-editor.org/info/rfc7015>>.

[RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, Ed., "TCP Extensions for High Performance", DOI 10.17487/RFC7323, RFC 7323, September 2014, <<https://www.rfc-editor.org/info/rfc7323>>.

[RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, RFC 7679, DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>.

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

Authors' Addresses

Thomas Graf
Swisscom
Binzring 17
CH-8045 Zurich
Switzerland

Email: thomas.graf@swisscom.com

Benoit Claise
Huawei

Email: benoit.claise@huawei.com

Alex Huang Feng
INSA-Lyon

Lyon
France

Email: alex.huang-feng@insa-lyon.fr