

Internet Engineering Task Force (IETF)
Request for Comments: 9951
Category: Standards Track
ISSN: 2070-1721

T. Graf
Swisscom
B. Claise
Huawei
A. Huang-Feng
INSA-Lyon
April 2026

Export of Delay Performance Metrics in IP Flow Information Export (IPFIX)

Abstract

This document specifies new IP Flow Information Export (IPFIX) Information Elements to export the On-Path delay at each Operations, Administration, and Maintenance (OAM) transit and decapsulating nodes. The On-Path delay is defined as the delay between the OAM header encapsulating node and each OAM header transit and OAM header decapsulating nodes. This delay measurement is computed by an On-Path Telemetry protocol and is exported by the IPFIX process.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9951>.

Copyright Notice

Copyright (c) 2026 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. Terminology
3. Solution
4. Performance Metrics
 - 4.1. IP One-Way Delay Hybrid Type I Performance Metrics
 - 4.1.1. Summary
 - 4.1.2. Description
 - 4.1.3. Reference
 - 4.1.4. Change Controller

- 4.1.5. Version of Registry Format
- 4.2. Metric Definition
 - 4.2.1. Reference Definition
 - 4.2.2. Fixed Parameters
- 4.3. Method of Measurement
 - 4.3.1. Reference Methods
 - 4.3.2. Packet Stream Generation
 - 4.3.3. Traffic Filtering (Observation) Details
 - 4.3.4. Sampling Distribution
 - 4.3.5. Runtime Parameters and Data Format
 - 4.3.6. Roles
- 4.4. Output
 - 4.4.1. Type
 - 4.4.2. Reference Definition
 - 4.4.3. Administrative Items
 - 4.4.4. Comments and Remarks
- 5. Use Cases
- 6. IANA Considerations
 - 6.1. Performance Metrics
 - 6.2. IPFIX Entities
 - 6.2.1. pathDelayMeanDeltaMicroseconds
 - 6.2.2. pathDelayMinDeltaMicroseconds
 - 6.2.3. pathDelayMaxDeltaMicroseconds
 - 6.2.4. pathDelaySumDeltaMicroseconds
- 7. Operational Considerations
 - 7.1. Time Accuracy
 - 7.2. Mean Delay
 - 7.3. Reduced-Size Encoding
 - 7.4. Measurement Interval
 - 7.5. In-Packet OAM Application
- 8. Security Considerations
- 9. References
 - 9.1. Normative References
 - 9.2. Informative References
- Appendix A. IPFIX Encoding Examples
 - A.1. Aggregated On-Path Delay Examples
 - A.1.1. Template Record and Data Set with Mean Delta
 - A.1.2. Template Record and Data Set with Sum Delta
- Acknowledgements
- Authors' Addresses

1. Introduction

Network operators usually maintain statistical views of delay across their networks to support diagnostics and performance analysis. These views assist in identifying the location, extent, and potential causes of abnormal delay affecting specific customer traffic or services. To achieve this, delay-related metrics need to be reported from devices covering both data and control planes. Further, in order to understand which customers are affected, delay-related metrics need to be reported in the context of the customer data plane. This correlation enables the detection of changes in forwarding paths, such as updated intermediate hops or interfaces, and of the resulting impact on delay experienced by customer traffic.

Delay measurements in the network are computed using an On-Path Telemetry protocol, which inserts metadata into the data-plane packet when entering the monitored domain [RFC9232]. To compute delay measurements, the On-Path Telemetry protocol inserts a timestamp reference when entering the OAM encapsulating node. Implementation examples are In situ Operations, Administration, and Maintenance (IOAM) [RFC9197] or Enhanced Alternate Marking [ENH-ALT-MARKING].

Two modes of On-Path Telemetry are generally recognized: passport mode, in which only the OAM header decapsulating node of the OAM domain reports metrics; and postcard mode, in which OAM header

transit nodes also export On-Path Telemetry data. Both modes enable exposure of per-hop performance metrics, including delay accumulation. The approach defined in this document is primarily applicable to postcard mode.

To enable the export of the delay-related metrics via IPFIX [RFC7011], this document defines four new IPFIX Information Elements (IEs), exposing the On-Path delay on OAM header transit and decapsulating nodes, following the principles of postcard mode.

This enables the computation of delay metrics (minimum, maximum, and mean) directly on the OAM header transit and decapsulating node, allowing aggregation within the Flow Record.

As these IEs represent performance metrics, they are also registered in the IANA "Performance Metrics Registry" [IANA-PERF-METRIC] in accordance with [RFC8911].

2. Terminology

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.

This document defines the following terms:

OAM Header Encapsulating Node:

Receives the IP packets, encapsulates the packets with an OAM header, and adds the timestamp into the OAM header.

OAM Header Transit Node:

Receives the IP packets, then measures the delay between the timestamp in the packet and the timestamp when the packet was received.

OAM Header Decapsulating Node:

Receives the IP packets, computes the delay between the timestamp in the packet and the timestamp when the packet was received, and removes the OAM header from the packet.

This document makes use of the terms defined in [RFC7011], [RFC8911] and [RFC7799].

The following terms are used as defined in [RFC7011]:

- * Collector
- * Exporter
- * Flow
- * Flow Record
- * IPFIX
- * IPFIX Information Elements (IEs)
- * Observation Point

The following terms are used as defined in [RFC8911]:

- * Performance Metric
- * Performance Metrics Registry

* Registered Performance Metric

The following term is used as defined in Section 3.8 of [RFC7799]:

* Hybrid Type I

3. Solution

In line with the guidelines for Registered Performance Metric requesters and reviewers [RFC8911], each metric is specified with its required characteristics (e.g., Identifier, Name, URI, Status, Requester, Revision, Description) to ensure comparability of measurement results across implementations and environments. These characteristics are registered in the IANA "Performance Metrics Registry" [IANA-PERF-METRIC]. Metric naming follows the "MetricType_Method_SubTypeMethod... Spec_Units_Output" convention defined in Section 7.1.2 of [RFC8911].

This document defines the following performance metrics and IPFIX Information Elements:

Performance Metric	IPFIX Information Element
OWDelay_HybridType1_I P_RFC9951_Seconds_Mean (27)	pathDelayMeanDeltaMicroseconds (530)
OWDelay_HybridType1_I P_RFC9951_Seconds_Min (28)	pathDelayMinDeltaMicroseconds (531)
OWDelay_HybridType1_I P_RFC9951_Seconds_Max (29)	pathDelayMaxDeltaMicroseconds (532)
OWDelay_HybridType1_I P_RFC9951_Seconds_Sum (30)	pathDelaySumDeltaMicroseconds (533)

Table 1: Mapping Between IPFIX IEs and Performance Metrics

Assuming time synchronization on devices, the delay is measured by calculating the difference between the timestamp imposed with On-Path Telemetry in the packet at an OAM header encapsulating node and the timestamp exported in the IPFIX Flow Record from the OAM header transit and OAM header decapsulating nodes. The lowest, highest, mean, and the sum of measured path delay can be exported, thanks to the different IPFIX IE specifications.

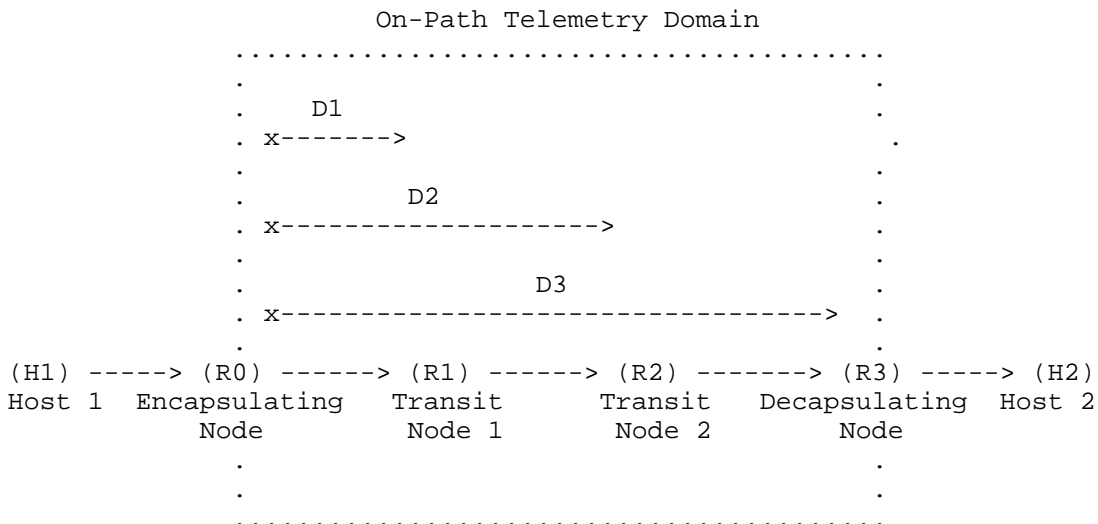


Figure 1: Delay Use Case: Packets Flow from Host 1 to Host 2

In the use case shown in Figure 1, using On-Path Telemetry to export the delay metrics, the node R1 exports the delay D1, the node R2 exports the delay D2, and the OAM header decapsulating node R3 exports the total delay D3 for the same flow using IPFIX.

This solution enables the computation of delay metrics (minimum, maximum, and mean) directly on the OAM header transit and decapsulating node, allowing aggregation within the Flow Record. This reduces both export bandwidth and processing requirements on the Collector. To compute these metrics locally, the Exporter's Metering Process must perform per-packet caching and processing, particularly when computing mean delay under Flow Aggregation [RFC7015]. A less computationally intensive alternative is to export the sum of delays, allowing the Collector to compute the mean via a simple division using the packet count.

In contrast, if no delay processing occurs on the OAM header transit or decapsulating node, each packet must be exported as an individual Flow Record, including timestamp information, as specified in [IPFIX-ALT-MARK]. The Collector must then compute the delay metrics and reconstruct the aggregated Flow Record accordingly.

4. Performance Metrics

This section defines four new performance metrics following the template defined in Section 11 of [RFC8911].

4.1. IP One-Way Delay Hybrid Type I Performance Metrics

This section specifies four performance metrics for the Hybrid Type I assessment of IP One-Way Delay; they have been registered in the IANA "Performance Metrics Registry" [IANA-PERF-METRIC].

All column entries besides the Identifier, Name, URI, Description, Reference Description (Output only) categories are the same; thus, this section defines four closely related performance metrics. As a result, IANA has assigned corresponding URIs to each of the four registered performance metrics.

4.1.1. Summary

This category includes multiple indexes of the registered performance metrics: the Identifier and Metric Name.

4.1.1.1. ID (Identifier)

IANA has allocated the numeric Identifiers 27, 28, 29, and 30 for the four Named Metric Entries in the following section.

4.1.1.2. Name

27: OWDelay_HybridType1_IP_RFC9951_Seconds_Mean
28: OWDelay_HybridType1_IP_RFC9951_Seconds_Min
29: OWDelay_HybridType1_IP_RFC9951_Seconds_Max
30: OWDelay_HybridType1_IP_RFC9951_Seconds_Sum

4.1.1.3. URI

URI:
<<https://www.iana.org/assignments/performance-metrics/>

OWDelay_HybridType1_IP_RFC9951_Seconds_Mean>

URI:

<[https://www.iana.org/assignments/performance-metrics/
OWDelay_HybridType1_IP_RFC9951_Seconds_Min](https://www.iana.org/assignments/performance-metrics/OWDelay_HybridType1_IP_RFC9951_Seconds_Min)>

URI:

<[https://www.iana.org/assignments/performance-metrics/
OWDelay_HybridType1_IP_RFC9951_Seconds_Max](https://www.iana.org/assignments/performance-metrics/OWDelay_HybridType1_IP_RFC9951_Seconds_Max)>

URI:

<[https://www.iana.org/assignments/performance-metrics/
OWDelay_HybridType1_IP_RFC9951_Seconds_Sum](https://www.iana.org/assignments/performance-metrics/OWDelay_HybridType1_IP_RFC9951_Seconds_Sum)>

4.1.2. Description

OWDelay_HybridType1_IP_RFC9951_Seconds_Mean:

This metric assesses the mean of one-way delays of all successfully forwarded IP packets constituting a single Flow. The measurement of one-way delay is based on a single Observation Point [RFC7011] somewhere in the network.

OWDelay_HybridType1_IP_RFC9951_Seconds_Min:

This metric assesses the minimum of one-way delays of all successfully forwarded IP packets constituting a single Flow. The measurement of one-way delay is based on a single Observation Point [RFC7011] somewhere in the network.

OWDelay_HybridType1_IP_RFC9951_Seconds_Max:

This metric assesses the maximum of one-way delays of all successfully forwarded IP packets constituting a single Flow. The measurement of one-way delay is based on a single Observation Point [RFC7011] somewhere in the network.

OWDelay_HybridType1_IP_RFC9951_Seconds_Sum:

This metric assesses the sum of one-way delays of all successfully forwarded IP packets constituting a single Flow. The measurement of one-way delay is based on a single Observation Point [RFC7011] somewhere in the network.

4.1.3. Reference

RFC 9951

4.1.4. Change Controller

IETF

4.1.5. Version of Registry Format

1.0

4.2. 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".

4.2.1. Reference Definition

See [RFC6049] and [RFC7679] in the Normative References (Section 9.1).

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 11 of [RFC2330].

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

This document specifies how to export the performance metric using IPFIX.

4.2.2. Fixed Parameters

None

4.3. 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.

4.3.1. 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 Observation Point [RFC7011].

4.3.2. Packet Stream Generation

This is the time when the packet is being received at the OAM header encapsulating node. The timestamp format depends on the On-Path Telemetry implementation. For IOAM, Section 4.4.1 of [RFC9197] describes the supported timestamps. Sections 4.4.2.3 and 4.4.2.4 of [RFC9197] describe where the timestamp is being inserted. For the Enhanced Alternate Marking Method, Section 2 of [ENH-ALT-MARKING] and Section 3.2 of [RFC9947] define the supported timestamp encodings and granularity.

4.3.3. Traffic Filtering (Observation) Details

Runtime Parameters (in the following sections) may be used for Traffic Filtering.

4.3.4. Sampling Distribution

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

4.3.5. Runtime Parameters and Data Format

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

The Hybrid Type I metering parameters must be reported to provide the complete measurement context. As an example, if the IPFIX Metering Process is used, then the IPFIX Metering Process parameters (IPFIX Template Record, potential traffic filters, and potential sampling method and parameters) that generate the Flow Records must be reported to provide the complete measurement context. At a minimum, the following fields are required:

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 [RFC9911]).

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 [RFC9911]).

T0: T time, the start of a measurement interval (format "date/time" as specified in Section 5.6 of [RFC3339]; see also "date-and-time" in Section 3 of [RFC9911]). The UTC Time Zone is required by Section 6.1 of [RFC2330]. When T0 is "all-zeros", a start time is unspecified, and Tf is to be interpreted as the duration of the measurement interval. The start time is controlled through other means.

Tf: A time, the end of a measurement interval (format "date/time" as specified in Section 5.6 of [RFC3339]; see also "date-and-time" in Section 3 of [RFC9911]). The UTC Time Zone is required by Section 6.1 of [RFC2330]. When T0 is "all-zeros", an ending time and date are ignored, and Tf is interpreted as the duration of the measurement interval.

4.3.6. Roles

Host A: Launches an IP packet to start the Flow.

Host B: Receives the IP packet to start the Flow.

OAM Header Encapsulating Node: Receives the IP packets, encapsulates the packets with the OAM header, and adds the timestamp into the OAM header.

OAM Header Transit Node: Receives the IP packets, measures the delay between the timestamp in the packet and the timestamp when the packet was received.

OAM Header Decapsulating Node: Receives the IP packets, computes the delay between the timestamp in the packet and the timestamp when the packet was received, and removes the OAM header from the packet.

4.4. Output

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

4.4.1. Type

OWDelay Types are discussed in the subsections below.

4.4.2. Reference Definition

For all output types:

OWDelay_HybridType1_IP: The one-way delay of one IP packet is a singleton.

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

4.4.2.1. OWDelay_HybridType1_IP_RFC9951_Seconds_Mean

Similar to Section 7.4.2.2 of [RFC8912], 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 analytic 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 microseconds, as a positive value of type decimal64 with fraction digits = 9 (similar to decimal64 in YANG, Section 9.3 of [RFC6020]) with a resolution of 0.001 microseconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per Section 6 of [RFC5905].

4.4.2.2. OWDelay_HybridType1_IP_RFC9951_Seconds_Min

Similar to Section 7.4.2.3 of [RFC8912], 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 analytic 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 microseconds, as a positive value of type decimal64 with fraction digits = 9 (similar to decimal64 in YANG, Section 9.3 of [RFC6020]) with a resolution of 0.001 microseconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per Section 6 of [RFC5905].

4.4.2.3. OWDelay_HybridType1_IP_RFC9951_Seconds_Max

Similar to Section 7.4.2.4 of [RFC8912], 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 analytic 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

where all packets n = 1 through N have finite singleton delays.

Max: The time value of the result is expressed in units of microseconds, as a positive value of type decimal64 with fraction digits = 9 (similar to decimal64 in YANG, Section 9.3 of [RFC6020]) with a resolution of 0.001 microseconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per Section 6 of [RFC5905].

4.4.2.4. OWDelay_HybridType1_IP_RFC9951_Seconds_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 analytic 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 microseconds, as a positive value of type decimal64 with fraction digits = 9 (similar to decimal64 in YANG, Section 9.3 of [RFC6020]) with a resolution of 0.001 microseconds (1.0 ns), and with lossless conversion to/from the 64-bit NTP timestamp as per Section 6 of [RFC5905].

4.4.2.5. Metric Units

- * Mean
- * Min
- * Max
- * Sum

The one-way delay of the IP Flow singleton is expressed in microseconds.

4.4.2.6. Calibration

A clock synchronization between the nodes of the monitored OAM domain is needed to compute representative delay measurements at the OAM header transit and decapsulating nodes. NTP, as defined in [RFC5905], can be used for synchronizing the clocks of the monitored nodes.

4.4.3. Administrative Items

4.4.3.1. Status

Current

4.4.3.2. Requester

RFC 9951

4.4.3.3. Revision

1.0

4.4.3.4. Revision Date

2026-04-02

4.4.4. Comments and Remarks

None

5. Use Cases

The measured On-Path delay can be aggregated with Flow Aggregation as defined in [RFC7015] across the following device and control-plane dimensions [IANA-IPFIX] to determine:

- * With node id and egressInterface(14), 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(15) or ipNextHopIPv6Address(62), the forwarding path to which next-hop IP contributed to how much delay.
- * With mplsTopLabelIPv4Address(47) or destinationIPv6Address and srhActiveSegmentIPv6(495), the forwarding path to which MPLS top-label IPv4 address or IPv6 destination address and Segment Routing over IPv6 (SRv6) active segment contributed to how much delay.
- * BGP communities [RFC1997] 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(8), or equivalent IPFIX IEs for IPv6, the forwarding path delay on each node from each IPv4 source address to a specific application in the network.

Let us consider Figure 1 as a topology example. Table 2 shows the aggregated delay per each node, ingressInterface(10), egressInterface(14), destinationIPv6Address(28), and srhActiveSegmentIPv6(495) measured at ingress.

ingress Interface	egress Interface	Node	destination IPv6Address	srhActive SegmentIPv6	Path Delay
271	276	R0			0 s
301	312	R1	2001:db8::1	2001:db8::3	22 s
22	27	R2	2001:db8::2	2001:db8::3	42 s
852	854	R3	2001:db8::3	2001:db8::3	122 s

Table 2: Example Table of Measured Delay at Ingress, Ascending by Delay

6. IANA Considerations

6.1. Performance Metrics

IANA has add four new performance metrics in the "Performance Metrics Registry" [RFC8911] with the four templates defined in Section 3.

6.2. IPFIX Entities

IANA has registered new IPFIX IEs (see Table 3) in the "IPFIX Information Elements" registry in the "IP Flow Information Export (IPFIX) Entities" registry group [IANA-IPFIX] and assigned the following code points.

=====

ElementID	Name
530	pathDelayMeanDeltaMicroseconds
531	pathDelayMinDeltaMicroseconds
532	pathDelayMaxDeltaMicroseconds
533	pathDelaySumDeltaMicroseconds

Table 3: New IPFIX IEs in the "IPFIX Information Elements" Registry

6.2.1. pathDelayMeanDeltaMicroseconds

Name: pathDelayMeanDeltaMicroseconds

ElementID: 530

Description: This Information Element identifies the mean path delay of all packets in the Flow, in microseconds, between an OAM header encapsulating node and the local node with the OAM domain (either an OAM header transit node or an OAM header decapsulating node), according to OWDelay_HybridType1_IP_RFC9951_Seconds_Mean in the IANA "Performance Metrics Registry".

Abstract Data Type: unsigned32

Data Type Semantics: deltaCounter

Reference: RFC 9951

Additional Information: OWDelay_HybridType1_IP_RFC9951_Seconds_Mean in the IANA "Performance Metrics Registry".

6.2.2. pathDelayMinDeltaMicroseconds

Name: pathDelayMinDeltaMicroseconds

ElementID: 531

Description: This Information Element identifies the lowest path delay of all packets in the Flow, in microseconds, between an OAM header encapsulating node and the local node with the OAM domain (either an OAM header transit node or an OAM header decapsulating node), according to the OWDelay_HybridType1_IP_RFC9951_Seconds_Min in the IANA "Performance Metrics Registry".

Abstract Data Type: unsigned32

Data Type Semantics: deltaCounter

Reference: RFC 9951

Additional Information: OWDelay_HybridType1_IP_RFC9951_Seconds_Min in the IANA "Performance Metrics Registry".

6.2.3. pathDelayMaxDeltaMicroseconds

Name: pathDelayMaxDeltaMicroseconds

ElementID: 532

Description: This Information Element identifies the highest path delay of all packets in the Flow, in microseconds, between an OAM

header encapsulating node and the local node with the OAM domain (either an OAM header transit node or an OAM header decapsulating node), according to `OWDelay_HybridType1_IP_RFC9951_Seconds_Max` in the IANA "Performance Metrics Registry".

Abstract Data Type: `unsigned32`

Data Type Semantics: `deltaCounter`

Reference: RFC 9951

Additional Information: `OWDelay_HybridType1_IP_RFC9951_Seconds_Max` in the IANA "Performance Metrics Registry".

6.2.4. `pathDelaySumDeltaMicroseconds`

Name: `pathDelaySumDeltaMicroseconds`

ElementID: 533

Description: This Information Element identifies the sum of the path delay of all packets in the Flow, in microseconds, between an OAM header encapsulating node and the local node with the OAM domain (either an OAM header transit node or an OAM header decapsulating node), according to `OWDelay_HybridType1_IP_RFC9951_Seconds_Sum` in the IANA "Performance Metrics Registry".

Abstract Data Type: `unsigned64`

Data Type Semantics: `deltaCounter`

Reference: RFC 9951

Additional Information: `OWDelay_HybridType1_IP_RFC9951_Seconds_Sum` in the IANA "Performance Metrics Registry".

7. Operational Considerations

7.1. Time Accuracy

In terms of clock precision, the same recommendation as defined in Section 4.5 of [RFC5153] for IPFIX applies to this document as well.

7.2. Mean Delay

The mean (average) path delay can be calculated by dividing the `pathDelaySumDeltaMicroseconds(533)` by the `packetDeltaCount(2)` at the IPFIX data collection at the collection time instead of the IPFIX Exporter at the export time.

7.3. Reduced-Size Encoding

`Unsigned64` has been chosen as the type for `pathDelaySumDeltaMicroseconds` to support cases with large delay numbers and where many packets are being counted. As an example, a specific Flow Record with path delay of 100 milliseconds cannot observe more than 42949 packets without overflowing the `unsigned32` counter. The procedure described in Section 6.2 of [RFC7011] may be applied to reduce network bandwidth between the IPFIX Exporter and Collector if `unsigned32` would be large enough without wrapping around.

7.4. Measurement Interval

The delay metrics are computed for the Flow Record lifetime by comparing the OAM timestamps in each received packet with the

timestamp when they were received. For a long-running Flow, the IPFIX Metering Process might miss the temporal distribution of the delay (for example, a longer delay only at the beginning of the Flow). If this is an operational problem, the IPFIX Metering Process might be configured with a smaller expiration timeout (see "Flow Expiration", Section 5.1.1 of [RFC5470]).

7.5. In-Packet OAM Application

Multiple methods can be used to compute the delay performance metrics defined in this document. Some examples of such methods are IOAM [RFC9197] and Enhanced Alternate Marking [ENH-ALT-MARKING].

For IOAM, these performance metrics can be computed using the Edge-to-Edge and the Direct Exporting Option-Type.

IOAM Edge-to-Edge Option-Type, as described in Section 4.6 of [RFC9197], can use bits 2 and 3. In this case, timestamps are encoded as defined in Sections 4.4.2.3 and 4.4.2.4 of [RFC9197]. This timestamp can be used to compute the delay between an OAM header encapsulating node and the decapsulating node.

The IOAM Direct Exporting Option-Type, as described in [RFC9326], can use the Extension-Flag defined in [IOAM-DEX] to insert a timestamp in the OAM header encapsulating node. The timestamp is encoded as defined in Sections 4.4.2.3 and 4.4.2.4 of [RFC9197]. This timestamp can be used to compute the delay between the inserted timestamp and the OAM header transit and decapsulating node.

For the Enhanced Alternate Marking Method, Section 2 of [ENH-ALT-MARKING] and Section 3.2 of [RFC9947] define that, within the metaInfo, a nanosecond timestamp can be encoded in an OAM header encapsulating node and be read at the OAM header intermediate and decapsulating nodes to calculate the On-path delay. [RFC9343] defines how this can be applied to the IPv6 extensions header, and [RFC9947] defines how this can be applied to the SRv6 Segment Routing Header [RFC8754].

Given that the delay measurements are computed with the timestamp introduced on the OAM header encapsulating node, regardless of the approach, implementations should document at which point of the forwarding plane this timestamp is introduced (e.g., the time at which the packet was received by the node, the time at which the packet was transmitted by the node, etc.). Based on this information, different actions can be taken.

8. Security Considerations

The IPFIX Information Elements introduced in this document do not directly introduce security issues. Rather, they define a set of performance metrics that may, for privacy or business issues, be considered sensitive information.

For example, exporting delay metrics may make attacks possible by the receiver of this information; otherwise, this would only be possible for direct observers of the reported Flows along the data path.

IPFIX collectors MUST ensure that IPFIX data originates from trusted sources. Accepting IPFIX data from unauthenticated sources could lead to data spoofing, policy misapplication, or denial of service.

Therefore, the underlying protocol used to exchange the information described here must apply appropriate procedures to guarantee the integrity and confidentiality of the exported information. These protocols are defined in separate documents; specifically, see the IPFIX security considerations in Section 11 of [RFC7011].

Implementations SHOULD also refer to [BCP195] for additional details.

9. References

9.1. Normative References

- [BCP195] Best Current Practice 195,
<<https://www.rfc-editor.org/info/bcp195>>.
At the time of writing, this BCP comprises the following:
- Moriarty, K. and S. Farrell, "Deprecating TLS 1.0 and TLS 1.1", BCP 195, RFC 8996, DOI 10.17487/RFC8996, March 2021, <<https://www.rfc-editor.org/info/rfc8996>>.
- Sheffer, Y., Saint-Andre, P., and T. Fossati, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 9325, DOI 10.17487/RFC9325, November 2022, <<https://www.rfc-editor.org/info/rfc9325>>.
- [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>>.
- [RFC3339] Klyne, G. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, DOI 10.17487/RFC3339, July 2002, <<https://www.rfc-editor.org/info/rfc3339>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, Ed., "TCP Extensions for High Performance", RFC 7323, DOI 10.17487/RFC7323, 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>>.
- [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>>.
- [RFC8911] Bagnulo, M., Claise, B., Eardley, P., Morton, A., and A. Akhter, "Registry for Performance Metrics", RFC 8911, DOI 10.17487/RFC8911, November 2021, <<https://www.rfc-editor.org/info/rfc8911>>.
- [RFC8912] Morton, A., Bagnulo, M., Eardley, P., and K. D'Souza, "Initial Performance Metrics Registry Entries", RFC 8912, DOI 10.17487/RFC8912, November 2021,

<<https://www.rfc-editor.org/info/rfc8912>>.

9.2. Informative References

[ENH-ALT-MARKING]

Zhou, T., Ed., Fioccola, G., Liu, Y., Cociglio, M., Pang, R., Xiong, L., Lee, S., and W. Li, "Enhanced Alternate Marking Method", Work in Progress, Internet-Draft, draft-zhou-ippm-enhanced-alternate-marking-18, 5 December 2025, <<https://datatracker.ietf.org/doc/html/draft-zhou-ippm-enhanced-alternate-marking-18>>.

[IANA-IPFIX]

IANA, "IP Flow Information Export (IPFIX) Entities", <<https://www.iana.org/assignments/ipfix>>.

[IANA-PERF-METRIC]

IANA, "Performance Metrics", <<https://www.iana.org/assignments/performance-metrics>>.

[IOAM-DEX] Huang Feng, A., Francois, P., Claise, B., and T. Graf, "Timestamp extension for In Situ Operations, Administration, and Maintenance (IOAM) Direct Export", Work in Progress, Internet-Draft, draft-ahuang-ippm-dex-timestamp-ext-00, 15 February 2023, <<https://datatracker.ietf.org/doc/html/draft-ahuang-ippm-dex-timestamp-ext-00>>.

[IPFIX-ALT-MARK]

Graf, T., Fioccola, G., Zhou, T., and Y. Zhu, "IP Flow Information Export (IPFIX) Alternate-Marking Information Elements", Work in Progress, Internet-Draft, draft-ietf-opsawg-ipfix-alt-mark-05, 27 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-opsawg-ipfix-alt-mark-05>>.

[RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<https://www.rfc-editor.org/info/rfc1997>>.

[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>>.

[RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, DOI 10.17487/RFC5470, March 2009, <<https://www.rfc-editor.org/info/rfc5470>>.

[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>>.

[RFC6703] Morton, A., Ramachandran, G., and G. Maguluri, "Reporting

IP Network Performance Metrics: Different Points of View",
RFC 6703, DOI 10.17487/RFC6703, August 2012,
<<https://www.rfc-editor.org/info/rfc6703>>.

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

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

[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/info/rfc9197>>.

[RFC9232] Song, H., Qin, F., Martinez-Julia, P., Ciavaglia, L., and A. Wang, "Network Telemetry Framework", RFC 9232, DOI 10.17487/RFC9232, May 2022, <<https://www.rfc-editor.org/info/rfc9232>>.

[RFC9326] Song, H., Gafni, B., Brockners, F., Bhandari, S., and T. Mizrahi, "In Situ Operations, Administration, and Maintenance (IOAM) Direct Exporting", RFC 9326, DOI 10.17487/RFC9326, November 2022, <<https://www.rfc-editor.org/info/rfc9326>>.

[RFC9343] Fioccola, G., Zhou, T., Cociglio, M., Qin, F., and R. Pang, "IPv6 Application of the Alternate-Marking Method", RFC 9343, DOI 10.17487/RFC9343, December 2022, <<https://www.rfc-editor.org/info/rfc9343>>.

[RFC9911] Schnwlder, J., Ed., "Common YANG Data Types", RFC 9911, DOI 10.17487/RFC9911, December 2025, <<https://www.rfc-editor.org/info/rfc9911>>.

[RFC9947] Fioccola, G., Zhou, T., Mishra, G., Wang, X., Zhang, G., and M. Cociglio, "Application of the Alternate-Marking Method to the Segment Routing Header", RFC 9947, DOI 10.17487/RFC9947, March 2026, <<https://www.rfc-editor.org/info/rfc9947>>.

Appendix A. IPFIX Encoding Examples

This appendix represents two different encodings for the newly introduced IEs. Let's take Figure 1 as a topology example. Table 4 shows the aggregated delay with ingressInterface, egressInterface, destinationIPv6Address, and srhActiveSegmentIPv6.

ingressInterface	271	
egressInterface	276	
destinationIPv6Address	2001:db8::3	
srhActiveSegmentIPv6	2001:db8::2	
packetDeltaCount	5	

pathDelayMeanDeltaMicroseconds	36 s	
pathDelayMinDeltaMicroseconds	22 s	
pathDelayMaxDeltaMicroseconds	74 s	

Table 4: Aggregated Delay with
egressInterface and srhActiveSegmentIPv6

A.1. Aggregated On-Path Delay Examples

A.1.1. Template Record and Data Set with Mean Delta

With encoding in Figure 2, the mean (average) path delay is calculated on the exporting node.

- * Ingress interface => ingressInterface
- * Egress interface => egressInterface
- * IPv6 destination address => destinationIPv6Address
- * Active SRv6 Segment => srhActiveSegmentIPv6
- * Packet Delta Count => packetDeltaCount
- * Minimum One-Way Delay => pathDelayMinDeltaMicroseconds (531)
- * Maximum One-Way Delay => pathDelayMaxDeltaMicroseconds (532)
- * Mean One-Way Delay => pathDelayMeanDeltaMicroseconds (530)

0	1	2	3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1			
SET ID = 2	Length = 40		
Template ID = 256	Field Count = 8		
0 ingressInterface = 10	Field Length = 4		
0 egressInterface = 14	Field Length = 4		
0 destinationIPv6Address = 28	Field Length = 16		
0 srhActiveSegmentIPv6 = 495	Field Length = 16		
0 packetDeltaCount = 5	Field Length = 4		
0 pathDelayMeanDelta.. = 530	Field Length = 4		
0 pathDelayMinDelta.. = 531	Field Length = 4		
0 pathDelayMaxDelta.. = 532	Field Length = 4		

Figure 2: Template Record for pathDelayMeanDeltaMicroseconds

The data set is represented as follows:

0	1	2	3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1			
SET ID = 256	Length = 60		

```

+-----+
| ingressInterface = 271 |
+-----+
| egressInterface = 276 |
+-----+
| destinationIPv6Address = |
| ... |
| ... |
| 2001:db8::2 |
+-----+
| srhActiveSegmentIPv6 = ... |
| ... |
| ... |
| 2001:db8::3 |
+-----+
| packetDeltaCount = 5 |
+-----+
| pathDelayMeanDeltaMicroseconds = 36 |
+-----+
| pathDelayMinDeltaMicroseconds = 22 |
+-----+
| pathDelayMaxDeltaMicroseconds = 74 |
+-----+

```

Figure 3: Data Set Encoding for pathDelayMeanDeltaMicroseconds

A.1.2. Template Record and Data Set with Sum Delta

With encoding in Figure 4, the mean (average) path delay is calculated on the IPFIX data collection.

- * Ingress interface => ingressInterface
- * Egress interface => egressInterface
- * IPv6 destination address => destinationIPv6Address
- * Active SRv6 Segment => srhActiveSegmentIPv6
- * Packet Delta Count => packetDeltaCount
- * Minimum One-Way Delay => pathDelayMinDeltaMicroseconds (531)
- * Maximum One-Way Delay => pathDelayMaxDeltaMicroseconds (532)
- * Sum of One-Way Delay => pathDelaySumDeltaMicroseconds (533)

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+
| SET ID = 2 | Length = 40 |
+-----+
| Template ID = 257 | Field Count = 8 |
+-----+
| 0 | ingressInterface = 10 | Field Length = 4 |
+-----+
| 0 | egressInterface = 14 | Field Length = 4 |
+-----+
| 0 | destinationIPv6Address = 28 | Field Length = 16 |
+-----+
| 0 | srhActiveSegmentIPv6 = 495 | Field Length = 16 |
+-----+
| 0 | packetDeltaCount = 5 | Field Length = 4 |
+-----+
| 0 | pathDelayMinDelta.. = 531 | Field Length = 4 |
+-----+

```

```

|0| pathDelayMaxDelta.. = 532 | Field Length = 4 |
+-----+-----+-----+-----+-----+-----+
|0| pathDelaySumDelta.. = 533 | Field Length = 8 |
+-----+-----+-----+-----+-----+-----+

```

Figure 4: Template Record for pathDelaySumDeltaMicroseconds.

The data set is represented as follows:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+
|                SET ID = 257                | Length = 64 |
+-----+-----+-----+-----+-----+-----+
|                ingressInterface = 271        |
+-----+-----+-----+-----+-----+-----+
|                egressInterface = 276         |
+-----+-----+-----+-----+-----+-----+
|                destinationIPv6Address =      |
|                ...                          |
|                ...                          |
|                2001:db8::2                  |
+-----+-----+-----+-----+-----+-----+
|                srhActiveSegmentIPv6 = ...    |
|                ...                          |
|                ...                          |
|                2001:db8::3                  |
+-----+-----+-----+-----+-----+-----+
|                packetDeltaCount = 5          |
+-----+-----+-----+-----+-----+-----+
|                pathDelayMinDeltaMicroseconds = 22 |
+-----+-----+-----+-----+-----+-----+
|                pathDelayMaxDeltaMicroseconds = 74 |
+-----+-----+-----+-----+-----+-----+
|                pathDelaySumDeltaMicroseconds = 180 |
|                ...                          |
+-----+-----+-----+-----+-----+-----+

```

Figure 5: Data Set Encoding for pathDelaySumDeltaMicroseconds

Acknowledgements

The authors would like to thank Al Morton (Rest in Peace, Al), Justin Iurman, Giuseppe Fioccola, Yannick Buchs, Menachem Dodge, Martin Duke, Behcet Sarikaya, Mahesh Jethanandani, Linda Dunbar, Deb Cooley, Mike Bishop, Tim Wicinski, Gunter Van de Velde, and ric Vyncke for their review and valuable comments. Special thanks to Paul Aitken (as IPFIX Designated Expert), Greg Mirsky (as IP Performance Metrics Designated Expert), and to Med Boucadair for their very detailed feedback.

Authors' Addresses

Thomas Graf
 Swisscom
 Binzring 17
 CH-8045 Zurich
 Switzerland
 Email: thomas.graf@swisscom.com

Benoit Claise
 Huawei
 Email: benoit@everything-ops.net

Alex Huang-Feng
INSA-Lyon
Lyon
France
Email: alex.huang-feng@insa-lyon.fr