ippm
Internet-Draft
Intended status: Informational
Expires: 15 August 2024

M. Spiegel
Barefoot Networks, an Intel company
F. Brockners
Cisco
S. Bhandari
Thoughtspot
R. Sivakolundu
Cisco

12 February 2024

In-situ OAM raw data export with IPFIX
draft-spiegel-ippm-ioam-rawexport-07

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 how In-situ Operations, Administration, and Maintenance (IOAM) information can be exported in raw, i.e. uninterpreted, format from network devices to systems, such as monitoring or analytics systems using IPFIX.

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 15 August 2024.

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. IOAM data fields are defined in [RFC9197]. This document discusses how In-situ Operations, Administration, and Maintenance (IOAM) information can be exported in raw format, i.e. uninterpreted format, from network devices to systems, such as monitoring or analytics systems using IPFIX [RFC7011].

"Raw export of IOAM data" refers to a mode of operation where a node exports the IOAM data as it is received in the packet. The exporting node neither interprets, aggregates nor reformats the IOAM data before it is exported. Raw export of IOAM data is to support an operational model where the processing and interpretation of IOAM data is decoupled from the operation of encapsulating/updating/ decapsulating IOAM data, which is also referred to as IOAM data-plane operation. The figure below shows the separation of concerns for IOAM export: Exporting IOAM data is performed by the "IOAM node" which performs IOAM data-plane operation, whereas the interpretation of IOAM data is performed by the IOAM data processing system. The separation of concerns is to off-load interpretation, aggregation and formatting of IOAM data from the node which performs data-plane operations. In other words, a node which is focused on data-plane operations, i.e. forwarding of packets and handling IOAM data will not be tasked to also interpret the IOAM data, but can leave this task to another system. Note that for scalability reasons, a single IOAM node could choose to export IOAM data to several IOAM data processing systems.

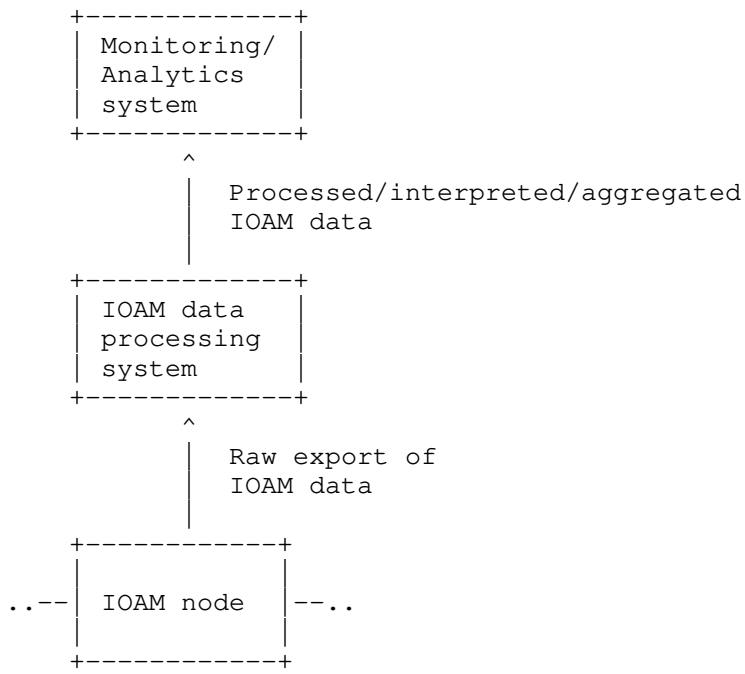

IOAM node: IOAM encapsulating, IOAM decapsulating or IOAM transit node.

IOAM data processing system: System that receives raw IOAM data and provides for formatting, aggregation and interpretation of the IOAM data.

Monitoring/Analytics system: System that receives telemetry and other operational information from a variety of sources and provides for correlation and interpretation of the data received.

Raw export of IOAM data is typically generated by network devices at the edges of the network. Deployment and use-case dependent, such as in case of direct export [RFC9326] or in cases where the operator is interested in dropped packets, raw export of IOAM data may be generated by IOAM transit nodes.

1.1. Requirements

Requirements for raw export of IOAM data:

* Export all IOAM information contained in a packet.

* Export a specific IOAM data type - Incremental Trace type, Preallocated Trace type, Proof of Transit type, Edge to Edge type, Direct Export type.

* Export IOAM trace data associated with a packet, even if that data was never included in a transmitted or received packet in the network, for example in case of direct export.

* Support coalescing of the IOAM data from multiple packets into a single raw export packet.

* Support export of additional parts of the packet, other than the IOAM data as part of the raw export. This could be parts of the packet header and/or parts of the packet payload. This additional information provides context to the IOAM data (e.g. to be used for flow identification) and is to enable the IOAM data processing system to perform further analysis on the received data.

* Report the reason why IOAM data was exported. The "reason for export" is to complement the IOAM data retrieved from the packet. For example, if a packet was dropped by a node due to congestion, it could be helpful to export the IOAM data of this dropped packet along with an indication that the packet that the IOAM data belongs to was dropped due to congestion.

1.2. Scope

This document discusses raw export of IOAM data using IPFIX.

The following is considered out of scope for this document:

* Protocols other than IPFIX for raw export of IOAM data.

* Interpretation or aggregation of IOAM data prior to exporting.

* Configuration of network devices so that they can determine when to generate IOAM reports, and what information to include in those reports.

* Events that trigger generation of IOAM reports.

* Selection of particular destinations within distributed telemetry monitoring systems, to which IOAM reports will be sent.

* Export format for flow statistics or processed/interpreted/ aggregated IOAM data.

2. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Abbreviations used in this document:

E2E: Edge to Edge

IOAM: In-situ Operations, Administration, and Maintenance

MTU: Maximum Transmit Unit

OAM: Operations, Administration, and Maintenance

POT: Proof of Transit

3. IPFIX for IOAM raw data export

IPFIX, being a generic export protocol, can export any Information Elements as long as they are described in the information model. The IPFIX protocol is well suited for and is defined as the protocol for exporting packet samples in [RFC5476].

IPFIX/PSAMP [RFC7011], [RFC5476] already define many of the information elements needed for exporting sections of packets needed for deriving context and raw IOAM data export. This document specifies extensions of the IPFIX information model for meeting the requirements in Section 1.1.

3.1. Key IPFIX information elements leveraged for IOAM raw data export

The existing IPFIX Information Elements that are required for IOAM raw data export are listed here. Their details are available in IANA’s IPFIX registry [IANA-IPFIX].

The existing IPFIX Information Elements used to carry the sections of the packets including IOAM data within it are as follows:

313 - ipHeaderPacketSection

315 - dataLinkFrameSection

The following Information Elements will be used to provide context to the ipHeaderPacketSection and dataLinkFrameSection as described in [IANA-IPFIX]:

408 - dataLinkFrameType

409 - sectionOffset

410 - sectionExportedOctets

The following Information Element will be used to provide forwarding status of the flow and any attached reasons.

89 - forwardingStatus

3.2. New IPFIX information elements leveraged for IOAM raw data export

IOAM data raw export using IPFIX requires a set of new information elements which are described in this section.

3.2.1. ioamReportFlags

Description:

This Information Element describes properties associated with an IOAM report.

The ioamReportFlags data type is an 8-bit field. The following bits are defined here:

Bit 0 Dropped Association - Dropped packet of interest.

Bit 1 Congested Queue Association - Indicates the presence of congestion on a monitored queue.

Bit 2 Tracked Flow Association - Matched a flow of interest.

Bit 3-7 Reserved

IANA is requested to create a new subregistry for IOAM Report Flags and fill it with the initial list from the description. New assignments for IOAM Encapsulation Types are administered by IANA through Expert Review [RFC5226] i.e., review by one of a group of experts designated by an IETF Area Director.

Abstract Data Type: unsigned8

Data Type Semantics: flags

ElementId: TBD1

Status: current

3.2.2. ioamEncapsulationType

Description:

This Information Element specifies the type of encapsulation to interpret ioamPreallocatedTraceData, ioamIncrementalTraceData, ioamE2EData, ioamPOTData, ioamDirectExportData.

The following ioamEncapsulationType values are defined here:

0 None : IOAM data follows format defined in [RFC9197]

1 GRE : IOAM data follows format defined in [I-D.weis-ippm-ioam-eth]

2 IPv6 : IOAM data follows format defined in [RFC9486]

3 VXLAN-GPE : IOAM data follows format defined in [I-D.brockners-ippm-ioam-vxlan-gpe]

4 GENEVE Option: IOAM data follows format defined in [I-D.brockners-ippm-ioam-geneve]

5 GENEVE Next Protocol: IOAM data follows format defined in [I-D.weis-ippm-ioam-eth]

6 NSH : IOAM data follows format defined in [RFC9452]

IANA is requested to create a new subregistry for IOAM Encapsulation Types and fill it with the initial list from the description. New assignments for IOAM Encapsulation Types are administered by IANA through Expert Review [RFC5226] i.e., review by one of a group of experts designated by an IETF Area Director.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

ElementId: TBD2

Status: current

3.2.3. ioamPreallocatedTraceData

Description:

This Information Element carries n octets of IOAM Preallocated Trace data defined in [RFC9197].

The format of the data is determined by the ioamEncapsulationType information element, if present. When the ioamEncapsulationType information element is present and has a value other than "None", and with sufficient length, this element may also report octets from subsequent headers and payload. If no ioamEncapsulationType information element is present, then the encapsulation type shall be assumed to be "None" and this information element only contains octets from the IOAM Preallocated Trace Option.

Abstract Data Type: octetArray

ElementId: TBD3

Status: current

3.2.4. ioamIncrementalTraceData

Description:

This Information Element carries n octets of IOAM Incremental Trace data defined in [RFC9197].

The format of the data is determined by the ioamEncapsulationType information element, if present. When the ioamEncapsulationType information element is present and has a value other than "None", and with sufficient length, this element may also report octets from subsequent headers and payload. If no ioamEncapsulationType information element is present, then the encapsulation type shall be assumed to be "None" and this information element only contains octets from the IOAM Incremental Trace Option.

Abstract Data Type: octetArray

ElementId: TBD4

Status: current

3.2.5. ioamE2EData

Description:

This Information Element carries n octets of IOAM E2E data defined in [RFC9197].

The format of the data is determined by the ioamEncapsulationType information element, if present. When the ioamEncapsulationType information element is present and has a value other than "None", and with sufficient length, this element may also report octets from

subsequent headers and payload. If no ioamEncapsulationType information element is present, then the encapsulation type shall be assumed to be "None" and this information element only contains octets from the IOAM Edge-to-Edge Option.

Abstract Data Type: octetArray

ElementId: TBD5

Status: current

3.2.6. ioamPOTData

Description:

This Information Element carries n octets of IOAM POT data defined in [RFC9197].

The format of the data is determined by the ioamEncapsulationType information element, if present. When the ioamEncapsulationType information element is present and has a value other than "None", and with sufficient length, this element may also report octets from subsequent headers and payload. If no ioamEncapsulationType information element is present, then the encapsulation type shall be assumed to be "None" and this information element only contains octets from the IOAM Proof of Transit Option.

Abstract Data Type: octetArray

ElementId: TBD6

Status: current

3.2.7. ioamDirectExportData

Description:

This Information Element carries n octets of IOAM Direct Export data defined in [RFC9326].

In addition to the fields from the IOAM Direct Export Option header in the packet, this information element includes all of the trace data from the exporting node, based on the IOAM-Trace-Type value. This data is appended inside ioamDirectExportData following the bit order of the IOAM-Trace-Type field, similar to the way that IOAM encapsulating nodes append trace data in Incremental Trace Option headers.

The format of the data is determined by the ioamEncapsulationType information element, if present. When the ioamEncapsulationType information element is present and has a value other than "None", and with sufficient length, this element may also report octets from subsequent headers and payload. If no ioamEncapsulationType information element is present, then the encapsulation type shall be assumed to be "None" and this information element only contains octets from the IOAM Direct Export Option plus the corresponding trace data.

Abstract Data Type: octetArray

ElementId: TBD7

Status: current

3.2.8. ipHeaderPacketSectionWithPadding

Description:

This Information Element carries a series of n octets from the IP header of a sampled packet, starting sectionOffset octets into the IP header.

However, if no sectionOffset field corresponding to this Information Element is present, then a sectionOffset of zero applies, and the octets MUST be from the start of the IP header.

With sufficient length, this element also reports octets from the IP payload. However, full packet capture of arbitrary packet streams is explicitly out of scope per the Security Considerations sections of [RFC5477] and [RFC2804].

When this Information Element has a fixed length, this MAY include padding octets that are used to fill out that fixed length.

When this information element has a variable length, the variable length MAY include up to 3 octets of padding, used to preserve 4-octet alignment of subequent Information Elements or subsequent records within the same set.

In either case of fixed or variable length, the amount of populated octets MAY be specified in the sectionExportedOctets field corresponding to this Information Element, in which case the remainder (if any) MUST be padding. If there is no sectionExportedOctets field corresponding to this Information Element, then all octets MUST be populated unless the total length of the IP packet is less than the fixed length of this Information Element, in which case the remainder MUST be padding.

Abstract Data Type: octetArray

ElementId: TBD8

Status: current

3.2.9. ethernetFrameSection

Description:

This Information Element carries a series of n octets from the IEEE 802.3 Ethernet frame of a sampled packet, starting after the preamble and start frame delimiter (SFD), plus sectionOffset octets into the frame if there is a sectionOffset field corresponding to this Information Element.

With sufficient length, this element also reports octets from the Ethernet payload. However, full packet capture of arbitrary packet streams is explicitly out of scope per the Security Considerations sections of [RFC5477] and [RFC2804].

When this Information Element has a fixed length, this MAY include padding octets that are used to fill out that fixed length.

When this information element has a variable length, the variable length MAY include up to 3 octets of padding, used to preserve 4-octet alignment of subequent Information Elements or subsequent records within the same set.

In either case of fixed or variable length, the amount of populated octets MAY be specified in the sectionExportedOctets field corresponding to this Information Element, in which case the remainder (if any) MUST be padding. If there is no sectionExportedOctets field corresponding to this Information Element, then all octets MUST be populated unless the total length of the Ethernet frame is less than the fixed length of this Information Element, in which case the remainder MUST be padding.

Abstract Data Type: octetArray

ElementId: TBD9

Status: current

4. Examples

This section shows a set of examples of how IOAM information along with other parts of the packet can be carried using IPFIX.

4.1. Fixed Length IP Packet

This example shows a fixed length IP packet. IOAM data is part of the ipHeaderPacketSection.

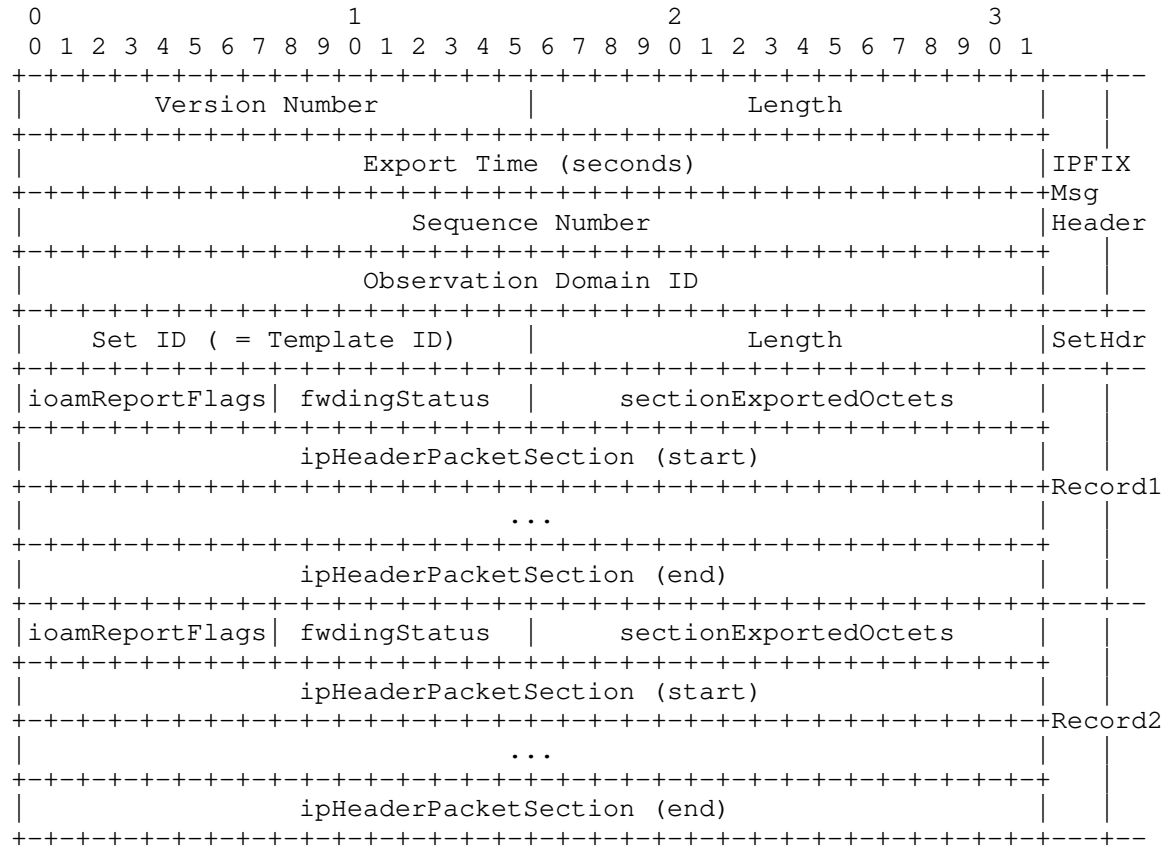

4.2. Variable Length IP Packet (length < 255)

This examples shows a variable length IP packet, with length < 255 bytes. IOAM data is part of the ipHeaderPacketSectionWithPadding.

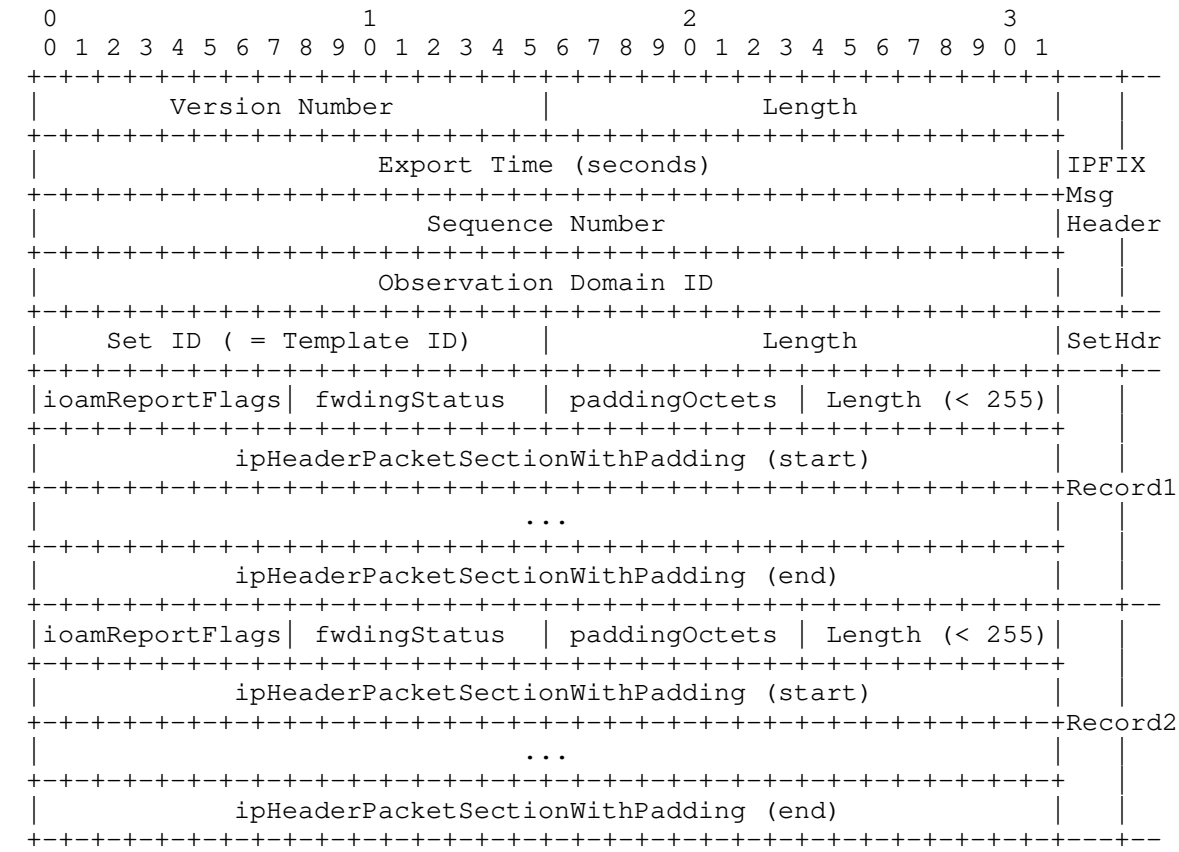

4.3. Variable Length IP Packet (length > 255)

This examples shows a variable length IP packet, with length > 255 bytes. IOAM data is part of the ipHeaderPacketSectionWithPadding.

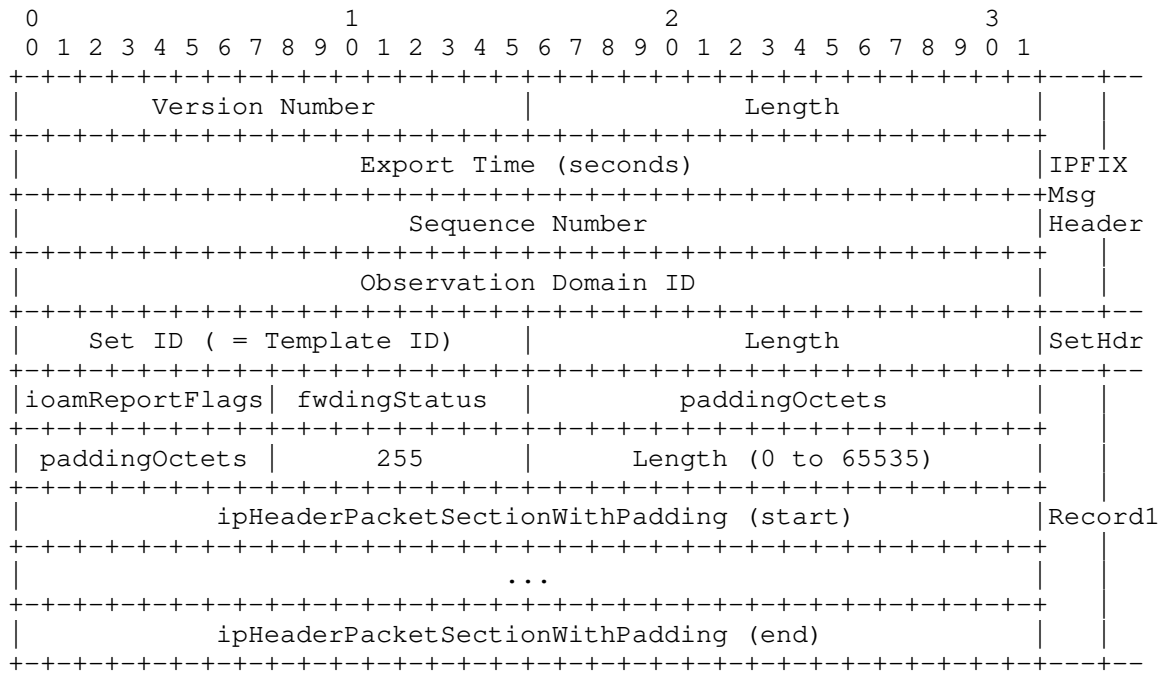

4.4. Variable Length ETHERNET Packet (length < 255)

This examples shows a variable length Ethernet packet, with length < 255 bytes. IOAM data is part of the ethernetFrameSection.

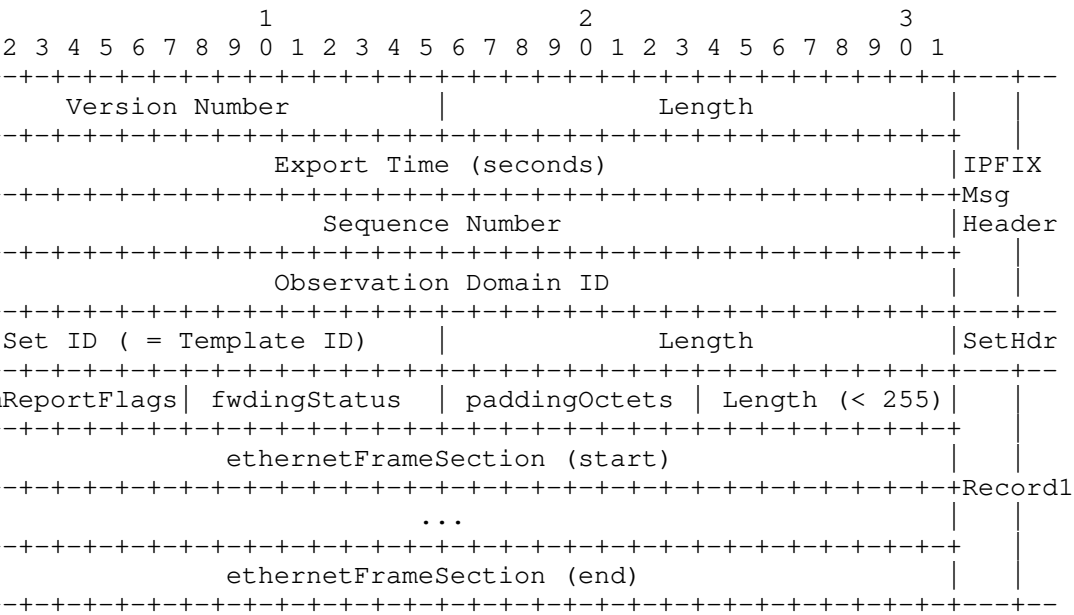

4.5. Variable Length IP Packet with Fixed Length IOAM Incremental Trace Data

This examples shows a variable length IP packet with length < 255 bytes and fixed length ioamIncrementalTraceData carried separately.

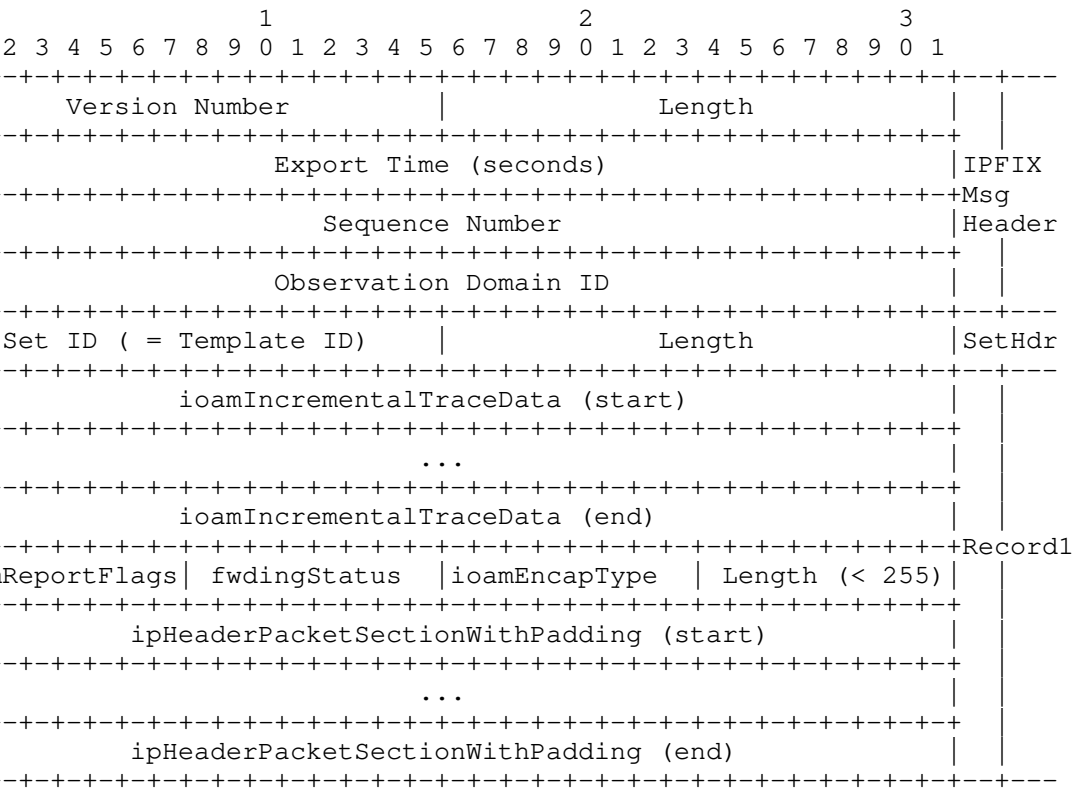

4.6. Variable Length IP Packet with Variable Length IOAM Incremental Trace Data

This examples shows a variable length IP packet with length < 255 bytes and variable length ioamIncrementalTraceData with length < 255 bytes carried separately.

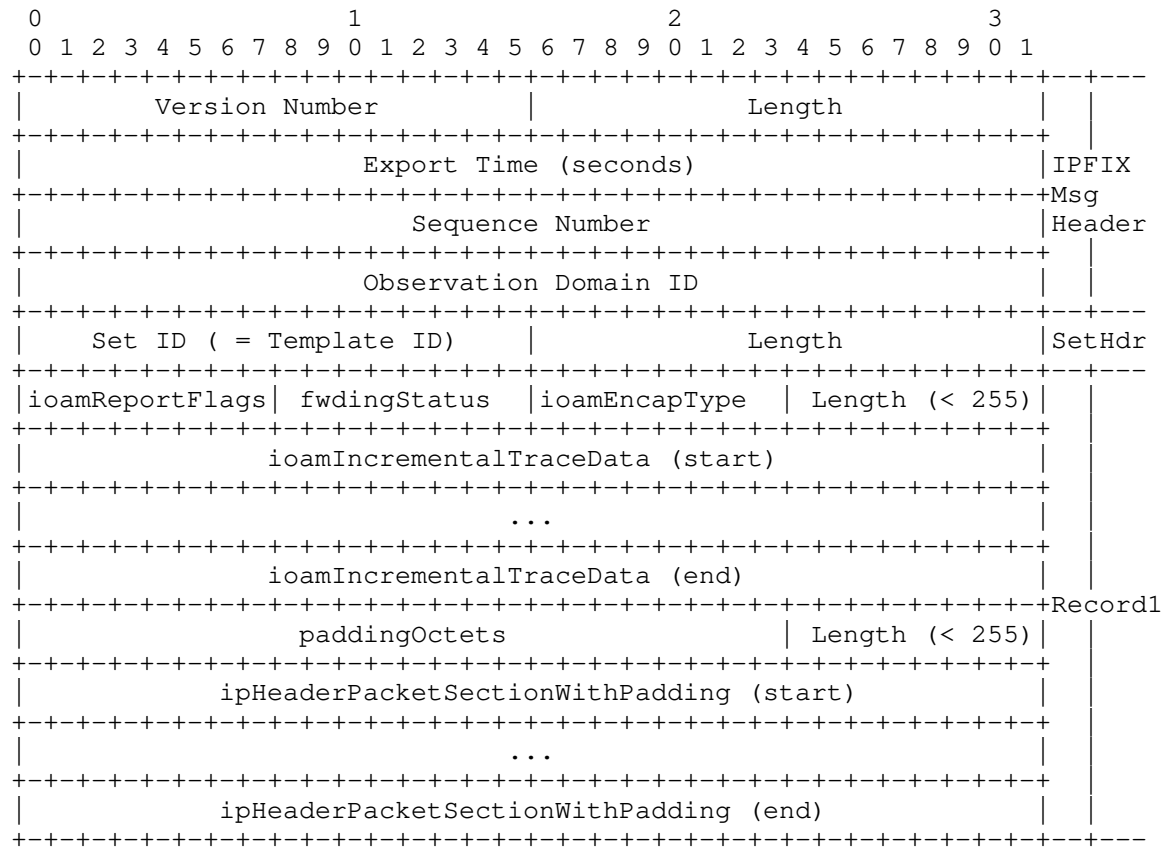

5. IANA Considerations

IANA is requested to allocate code points for the following Information Elements in [IANA-IPFIX]:

TBD1 ioamReportFlags

TBD2 ioamEncapsulationType

TBD3 ioamPreallocatedTraceData

TBD4 ioamIncrementalTraceData

TBD5 ioamE2EData

TBD6 ioamPOTData

TBD7 ioamDirectExportData

TBD8 ipHeaderPacketSectionWithPadding

TBD9 ethernetFrameSection

See Section 3.2 for further details.

IANA is requested to create subregistries for ioamReportFlags defined in Section 3.2.1 and ioamEncapsulationType defined in Section 3.2.2.

6. Manageability Considerations

Manageability considerations will be addressed in a later version of this document.

7. Security Considerations

Security considerations will be addressed in a later version of this document.

8. Acknowledgements

The authors would like to thank Barak Gafni, Tal Mizrahi, Jennifer Lemon, and Aviv Kfir for their thoughts and comments on raw IOAM data export.

9. References

9.1. 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>.

[RFC5476] Claise, B., Ed., Johnson, A., and J. Quittek, "Packet Sampling (PSAMP) Protocol Specifications", RFC 5476, DOI 10.17487/RFC5476, March 2009, <https://www.rfc-editor.org/info/rfc5476>.

[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>.

[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>.

9.2. Informative References

[I-D.brockners-ippm-ioam-geneve]
Brockners, F., Bhandari, S., Govindan, V. P., Pignataro, C., Nainar, N. K., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Lapukhov, P., Gafni, B., Kfir, A., and M. Spiegel, "Geneve encapsulation for In-situ OAM Data", Work in Progress, Internet-Draft, draft-brockners-ippm-ioam-geneve-05, 19 November 2020, <https://datatracker.ietf.org/doc/html/draft-brockners-ippm-ioam-geneve-05>.

[I-D.brockners-ippm-ioam-vxlan-gpe]
Brockners, F., Bhandari, S., Govindan, V. P., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., and M. Spiegel, "VXLAN-GPE Encapsulation for In-situ OAM Data", Work in Progress, Internet-Draft, draft-brockners-ippm-ioam-vxlan-gpe-04, 26 November 2023, <https://datatracker.ietf.org/doc/html/draft-brockners-ippm-ioam-vxlan-gpe-04>.

[I-D.weis-ippm-ioam-eth]
Weis, B., Brockners, F., Hill, C., Bhandari, S., Govindan, V. P., Pignataro, C., Nainar, N. K., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., and M. Spiegel, "EtherType Protocol Identification of In-situ OAM Data", Work in Progress, Internet-Draft, draft-weis-ippm-ioam-eth-05, 21 February 2022, <https://datatracker.ietf.org/doc/html/draft-weis-ippm-ioam-eth-05>.

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

[RFC2804] Internet Architecture Board and Internet Engineering Steering Group, "IETF Policy on Wiretapping", RFC 2804, DOI 10.17487/RFC2804, May 2000, <https://www.rfc-editor.org/info/rfc2804>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <https://www.rfc-editor.org/info/rfc5226>.

[RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", RFC 5477, DOI 10.17487/RFC5477, March 2009, <https://www.rfc-editor.org/info/rfc5477>.

[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>.

[RFC9452] Brockners, F., Ed. and S. Bhandari, Ed., "Network Service Header (NSH) Encapsulation for In Situ OAM (IOAM) Data", RFC 9452, DOI 10.17487/RFC9452, August 2023, <https://www.rfc-editor.org/info/rfc9452>.

[RFC9486] Bhandari, S., Ed. and F. Brockners, Ed., "IPv6 Options for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9486, DOI 10.17487/RFC9486, September 2023, <https://www.rfc-editor.org/info/rfc9486>.

Authors' Addresses

Mickey Spiegel
Barefoot Networks, an Intel company
4750 Patrick Henry Drive
Santa Clara, CA, 95054
United States of America
Email: mickey.spiegel@intel.com

Frank Brockners
Cisco Systems, Inc.
Hansaallee 249, 3rd Floor
40549 DUESSELDORF
Germany
Email: fbrockne@cisco.com

Shwetha Bhandari
Thoughtspot
3rd Floor, Indiqube Orion, 24th Main Rd, Garden Layout, HSR Layout
Bangalore, KARNATAKA 560 102
India
Email: shwetha.bhandari@thoughtspot.com

Ramesh Sivakolundu
Cisco Systems, Inc.
170 West Tasman Dr.
SAN JOSE, CA 95134,
United States of America
Email: sramesh@cisco.com