

ippm
Internet-Draft
Intended status: Informational
Expires: September 20, 2018

M. Spiegel
Barefoot Networks
F. Brockners
S. Bhandari
R. Sivakolundu
Cisco
March 19, 2018

In-situ OAM raw data export with IPFIX
draft-spiegel-ippm-ioam-rawexport-00

Abstract

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a path between two points in the network. This document discusses 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 <http://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 September 20, 2018.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements	4
1.2. Scope	5
2. Conventions	5
3. IPFIX for IOAM raw data export	6
3.1. Key IPFIX information elements leveraged for IOAM raw data export	6
3.2. New IPFIX information elements leveraged for IOAM raw data export	7
3.2.1. ioamReportFlags	7
3.2.2. ioamEncapsulationType	7
3.2.3. ioamPreallocatedTraceData	8
3.2.4. ioamIncrementalTraceData	9
3.2.5. ioamE2EData	9
3.2.6. ioamPOTData	10
4. Examples	10
4.1. Fixed Length IP Packet	10
4.2. Variable Length IP Packet (length < 255)	11
4.3. Variable Length IP Packet (length > 255)	12
4.4. Variable Length ETHERNET Packet (length < 255)	13
4.5. Variable Length IP Packet with Fixed Length IOAM Incremental Trace Data	14
4.6. Variable Length IP Packet with Variable Length IOAM Incremental Trace Data	15
5. IANA Considerations	16
6. Manageability Considerations	17
7. Security Considerations	17
8. Acknowledgements	17
9. References	17
9.1. Normative References	17
9.2. Informative References	18
Authors' Addresses	18

[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 [[I-D.ietf-ippm-ioam-data](#)]. This document discusses how In-situ Operations, Administration, and Maintenance (IOAM)

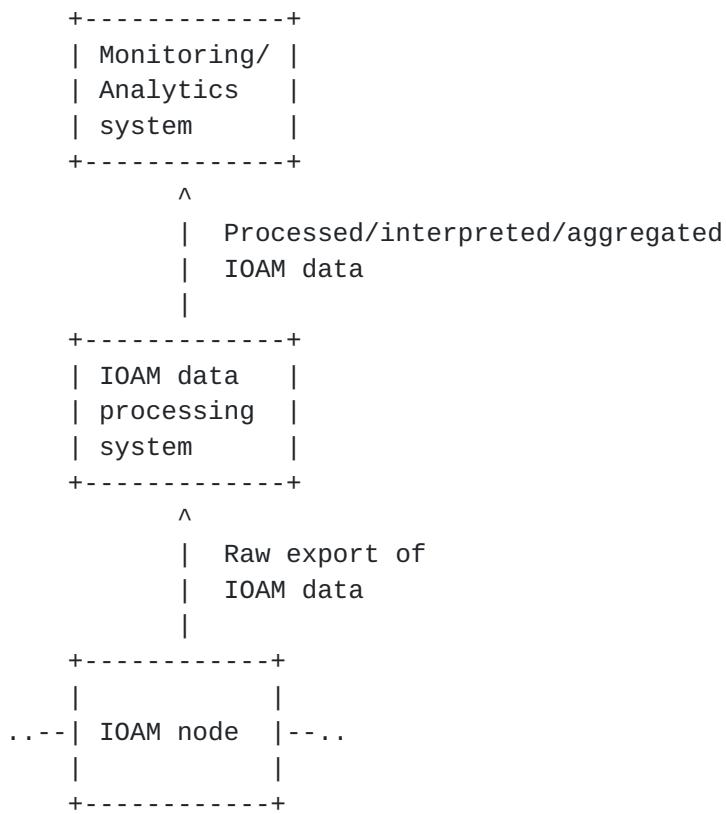
Spiegel, et al.

Expires September 20, 2018

[Page 2]

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 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:

- o Export all IOAM information contained in a packet.

Spiegel, et al.

Expires September 20, 2018

[Page 4]

- o Export a specific IOAM data type - Incremental Trace type, Preallocated Trace type, Proof of Transit type, Edge to Edge type.
- o Support coalescing of the IOAM data from multiple packets into a single raw export packet.
- o 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.
- o 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:

- o Protocols other than IPFIX for raw export of IOAM data.
- o Interpretation or aggregation of IOAM data prior to exporting.
- o Configuration of network devices so that they can determine when to generate IOAM reports, and what information to include in those reports.
- o Events that trigger generation of IOAM reports.
- o Selection of particular destinations within distributed telemetry monitoring systems, to which IOAM reports will be sent.
- o 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](#)].

Spiegel, et al.

Expires September 20, 2018

[Page 5]

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. IPFIX protocol is well suited for and is defined as the protocol for exporting packet samples in [[RFC5476](#)].

IPFIX/PSAMP [[RFC7011](#)], [[RFC5476](#)] already define most of the information elements needed for exporting sections of packets needed for deriving context and raw IOAM data export. This document specifies required 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 elements used to carry the sections of the packets including IOAM data within it are as follows:

313 - ipHeaderPacket[Section](#)

[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 ioamPreallocatedTraceHeader, ioamIncrementalTraceHeader, ioamE2EHeader, ioamPOTHeader.

The following ioamEncapsulationType values are defined here:

- 0 None : IOAM data follows format defined in
[[I-D.ietf-ippm-ioam-data](#)]
- 1 GRE : IOAM data follows format defined in [[I-D.weis-ippm-ioam-gre](#)]
- 2 VXLAN-GPE : IOAM data follows format defined in
[[I-D.brockners-ippm-ioam-vxlan-gpe](#)]
- 3 GENEVE : IOAM data follows format defined in
[[I-D.brockners-ippm-ioam-geneve](#)]
- 4 NSH : IOAM data follows format defined in
[[I-D.brockners-sfc-ioam-nsh](#)]

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 [[I-D.ietf-ippm-ioam-data](#)].

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

Spiegel, et al.

Expires September 20, 2018

[Page 8]

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 [[I-D.ietf-ipmm-ioam-data](#)].

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 [[I-D.ietf-ipmm-ioam-data](#)].

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

Spiegel, et al.

Expires September 20, 2018

[Page 9]

ElementId: TBD5

Status: current

[**3.2.6. ioamPOTData**](#)

Description:

This Information Element carries n octets of IOAM POT data defined in [[I-D.ietf-ipmm-ioam-data](#)].

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

[**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 examples 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 ipHeaderPacketSection.

Spiegel, et al.

Expires September 20, 2018

[Page 11]

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 ipHeaderPacketSection.

Spiegel, et al.

Expires September 20, 2018

[Page 12]

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 dataLinkFrameSection.

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.

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
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
Version Number		Length	
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
Export Time (seconds)			IPFIX
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+Message			
Sequence Number			Header
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
Observation Domain ID			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
Set ID (= Template ID)	Length		SetHdr
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ioamIncrementalTraceData (start)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
...			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ioamIncrementalTraceData (end)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+Record1			
ioamReportFlags fwdingStatus ioamEncapType Length (< 255)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ipHeaderPacketSection (start)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
...			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ipHeaderPacketSection (end)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			

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.

Spiegel, et al.

Expires September 20, 2018

[Page 15]

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
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
Version Number	Length		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
Export Time (seconds)	IPFIX		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+Message			
Sequence Number	Header		
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
Observation Domain ID			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
Set ID (= Template ID)	Length	SetHdr	
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ioamReportFlags fwdingStatus ioamEncapType Length (< 255)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ioamIncrementalTraceData (start)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
...			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ioamIncrementalTraceData (end)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+Record1			
paddingOctets Length (< 255)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ipHeaderPacketSection (start)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
...			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			
ipHeaderPacketSection (end)			
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+			

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

Spiegel, et al.

Expires September 20, 2018

[Page 16]

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, John Lemon, and Aviv Kfir for their thoughts and comments on raw IOAM data export.

9. References

[9.1. Normative References](#)

[I-D.ietf-ippm-ioam-data]

Brockners, F., Bhandari, S., Pignataro, C., Gredler, H.,
Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov,
P., Chang, R., daniel.bernier@bell.ca, d., and J. Lemon,
"Data Fields for In-situ OAM", [draft-ietf-ippm-ioam-data-02](#) (work in progress), March 2018.

[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>>.

Spiegel, et al.

Expires September 20, 2018

[Page 17]

9.2. Informative References

[I-D.brockners-ippm-ioam-geneve]

Brockners, F., Bhandari, S., Govindan, V., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Lapukhov, P., Gafni, B., Kfir, A., and M. Spiegel, "Geneve encapsulation for In-situ OAM Data", [draft-brockners-ippm-ioam-geneve-00](#) (work in progress), March 2018.

[I-D.brockners-ippm-ioam-vxlan-gpe]

Brockners, F., Bhandari, S., Govindan, V., 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", [draft-brockners-ippm-ioam-vxlan-gpe-00](#) (work in progress), March 2018.

[I-D.brockners-sfc-ioam-nsh]

Brockners, F., Bhandari, S., Govindan, V., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Mozes, D., Lapukhov, P., and R. Chang, "NSH Encapsulation for In-situ OAM Data", [draft-brockners-sfc-ioam-nsh-01](#) (work in progress), March 2018.

[I-D.weis-ippm-ioam-gre]

Weis, B., Brockners, F., crhill@cisco.com, c., Bhandari, S., Govindan, V., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., and M. Spiegel, "GRE Encapsulation for In-situ OAM Data", [draft-weis-ippm-ioam-gre-00](#) (work in progress), March 2018.

[IANA-IPFIX]

"IP Flow Information Export (IPFIX) Entities",
[<https://www.iana.org/assignments/ipfix/ipv4.xhtml>](https://www.iana.org/assignments/ipfix/ipv4.xhtml).

[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>>.

Authors' Addresses

Spiegel, et al.

Expires September 20, 2018

[Page 18]

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

Email: mspiegel@barefootnetworks.com

Frank Brockner
Cisco Systems, Inc.
Hansaallee 249, 3rd Floor
DUESSELDORF, NORDRHEIN-WESTFALEN 40549
Germany

Email: frockne@cisco.com

Shwetha Bhandari
Cisco Systems, Inc.
Cessna Business Park, Sarjapura Marathalli Outer Ring Road
Bangalore, KARNATAKA 560 087
India

Email: shwethab@cisco.com

Ramesh Sivakolundu
Cisco Systems, Inc.
170 West Tasman Dr.
SAN JOSE, CA 95134
U.S.A.

Email: sramesh@cisco.com

