Network Working Group
Request for Comments: 5103
Category: Standards Track

B. Trammell
CERT/NetSA
E. Boschi
Hitachi Europe

January 2008

Bidirectional Flow Export Using IP Flow Information Export (IPFIX)

Status of This Memo

Abstract

This document describes an efficient method for exporting bidirectional flow (Biflow) information using the IP Flow Information Export (IPFIX) protocol, representing each Biflow using a single Flow Record.

Table of Contents

1. Introduction

Many flow analysis tasks benefit from association of the upstream and downstream flows of a bidirectional communication, e.g., separating answered and unanswered TCP requests, calculating round trip times, etc. Metering processes that are not part of an asymmetric routing infrastructure, especially those deployed at a single point through which bidirectional traffic flows, are well positioned to observe bidirectional flows (Biflows). In such topologies, the total resource requirements for Biflow assembly are often lower if the Biflows are assembled at the measurement interface as opposed to the Collector. The IPFIX Protocol requires only information model extensions to be complete as a solution for exporting Biflow data.

To that end, we propose a Biflow export method using a single Flow Record per Biflow in this document. We explore the semantics of bidirectional flow data in Section 4, "Biflow Semantics"; examine the various possibilities for determining the direction of Biflows in Section 5, "Direction Assignment"; then define the Biflow export method in Section 6, "Record Representation".

This export method requires additional Information Elements to represent data values for the reverse direction of each Biflow, and a single additional Information Element to represent direction assignment information, as described in Sections 6.1 through 6.3. The selection of this method was motivated by an exploration of other possible methods of Biflow export using IPFIX; however, these methods have important drawbacks, as discussed in Section 3, "Rationale and History".

1.1. IPFIX Documents Overview

"Specification of the IPFIX Protocol for the Exchange of IP Traffic Flow Information" [RFC5101] (informally, the IPFIX Protocol document) and its associated documents define the IPFIX Protocol, which provides network engineers and administrators with access to IP traffic flow information.

"Architecture for IP Flow Information Export" [IPFIX-ARCH] (the IPFIX Architecture document) defines the architecture for the export of measured IP flow information out of an IPFIX Exporting Process to an IPFIX Collecting Process, and the basic terminology used to describe the elements of this architecture, per the requirements defined in "Requirements for IP Flow Information Export" [RFC3917]. The IPFIX Protocol document [RFC5101] then covers the details of the method for transporting IPFIX Data Records and Templates via a congestion-aware transport protocol from an IPFIX Exporting Process to an IPFIX Collecting Process.

"Information Model for IP Flow Information Export" [RFC5102] (informally, the IPFIX Information Model document) describes the Information Elements used by IPFIX, including details on Information Element naming, numbering, and data type encoding. Finally, "IPFIX Applicability" [IPFIX-AS] describes the various applications of the IPFIX protocol and their use of information exported via IPFIX, and relates the IPFIX architecture to other measurement architectures and frameworks.

This document references the Protocol and Architecture documents for terminology, uses the IPFIX Protocol to define a bidirectional flow export method, and proposes additions to the information model defined in the IPFIX Information Model document.

2. Terminology

Capitalized terms used in this document that are defined in the Terminology section of the IPFIX Protocol document [RFC5101] are to be interpreted as defined there. The following additional terms are defined in terms of the IPFIX Protocol document terminology.

Directional Key Field: A Directional Key Field is a single field in a Flow Key as defined in the IPFIX Protocol document [RFC5101] that is specifically associated with a single endpoint of the Flow. sourceIPv4Address and destinationTransportPort are example Directional Key Fields.

Non-directional Key Field: A Non-directional Key Field is a single field within a Flow Key as defined in the IPFIX Protocol document [RFC5101] that is not specifically associated with either endpoint of the Flow. protocolIdentifier is an example Non-directional Key Field.

Uniflow (Unidirectional Flow): A Uniflow is a Flow as defined in the IPFIX Protocol document [RFC5101], restricted such that the Flow is composed only of packets sent from a single endpoint to another single endpoint.

Biflow (Bidirectional Flow): A Biflow is a Flow as defined in the IPFIX Protocol document [RFC5101], composed of packets sent in both directions between two endpoints. A Biflow is composed from two Uniflows such that:

1. the value of each Non-directional Key Field of each Uniflow is identical to its counterpart in the other, and

2. the value of each Directional Key Field of each Uniflow is identical to its reverse direction counterpart in the other.

A Biflow contains two non-key fields for each value it represents associated with a single direction or endpoint: one for the forward direction and one for the reverse direction, as defined below.

Biflow Source: The Biflow Source is the endpoint identified by the source Directional Key Fields in the Biflow.

Biflow Destination: The Biflow Destination is the endpoint identified by the destination Directional Key Fields in the Biflow.

forward direction (of a Biflow): The direction of a Biflow composed of packets sent by the Biflow Source. Values associated with the forward direction of a Biflow are represented using normal Information Elements. In other words, a Uniflow may be defined as a Biflow having only a forward direction.

reverse direction (of a Biflow): The direction of a Biflow composed of packets sent by the Biflow Destination. Values associated with the reverse direction of a Biflow are represented using Reverse Information Elements, as defined below.

Reverse Information Element: An Information Element defined as corresponding to a normal (or forward) Information Element, but associated with the reverse direction of a Biflow.

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 RFC 2119 [RFC2119].

3. Rationale and History

In selecting the Single Record Biflow export method described in this document as the recommendation for bidirectional flow export using IPFIX, we considered several other possible methods.

The first and most obvious would be simply to export Biflows as two Uniflows adjacent in the record stream; a Collecting Process could then reassemble them with minimal state requirements. However, this has the drawbacks that it is merely an informal arrangement the Collecting Process cannot rely upon, and that it is not bandwidth-efficient, duplicating the export of Flow Key data in each Uniflow record.

We then considered the method outlined in Reducing Redundancy in IPFIX and Packet Sampling (PSAMP) Reports [IPFIX-REDUCING] for reducing this bandwidth inefficiency. This would also formally link

the two Uniflows into a single construct, by exporting the Flow Key as Common Properties then exporting each direction's information as Specific Properties. However, it would do so at the expense of additional overhead to transmit the commonPropertiesId, and additional state management requirements at both the Collecting and Exporting Processes.

A proposal was made on the IPFIX mailing list to use the Multiple Information Element feature of the protocol to export forward and reverse counters using identical Information Elements in the same Flow Record. In this approach, the first instance of a counter would represent the forward direction, and the second instance of the same counter would represent the reverse. This had the disadvantage of conflicting with the presently defined semantics for these counters, and, as such, was abandoned.

4. Biflow Semantics

As stated in the Terminology section above, a Biflow is simply a Flow representing packets flowing in both directions between two endpoints on a network. There are compelling reasons to treat Biflows as single entities (as opposed to merely ad-hoc combinations of Uniflows) within IPFIX. First, as most application-layer network protocols are inherently bidirectional, a Biflow-based data model more accurately represents the behavior of the network, and enables easier application of flow data to answering interesting questions about network behavior. Second, exporting Biflow data can result in improved export efficiency by eliminating the duplication of Flow Key data in an IPFIX message stream, and improve collection efficiency by removing the burden of Biflow matching from the Collecting Process where possible.

Biflows are somewhat more semantically complicated than Uniflows. When handling Uniflows, the semantics of source and destination Information Elements are clearly defined by the semantics of the underlying packet header data: the source Information Elements represent the source header fields, and the destination Information Elements represent the destination header fields. When representing Biflows with single IPFIX Data Records, the definitions of source and destination must be chosen more carefully.

As in the Terminology section above, we define the Source of a Biflow to be that identified by the source Directional Key Field(s), and the Destination of the Biflow to be that identified by the destination Directional Key Field(s). Note that, for IANA-registered Information Elements, or those defined by the IPFIX Information Model [RFC5102], Directional Key Fields associated with the Biflow Source are represented by Information Elements whose names begin with "source",

and Directional Key Fields associated with the Biflow Destination are represented by Information Elements whose names begin with "destination"; it is recommended that enterprise-specific Information Elements follow these conventions, as well.

Methods for assignment of Source and Destination by the Metering and Exporting Processes are described in the following section.

As the Source and Destination of a Biflow are defined in terms of its Directional Keys, Biflow values are also split info forward and reverse directions. As in the Terminology section above, the forward direction of a Biflow is composed of packets sent by the Biflow Source, and the reverse direction of a Biflow is composed of packets sent by the Destination. In other words, the two directions of a Biflow may be roughly thought of as the two Uniflows that were matched to compose the Biflow. A Biflow record, then, contains each Flow Key record once, and both forward Information Elements and Reverse Information Elements for each non-key field. See Figure 1 for an illustration of the composition of Biflows from Uniflows.

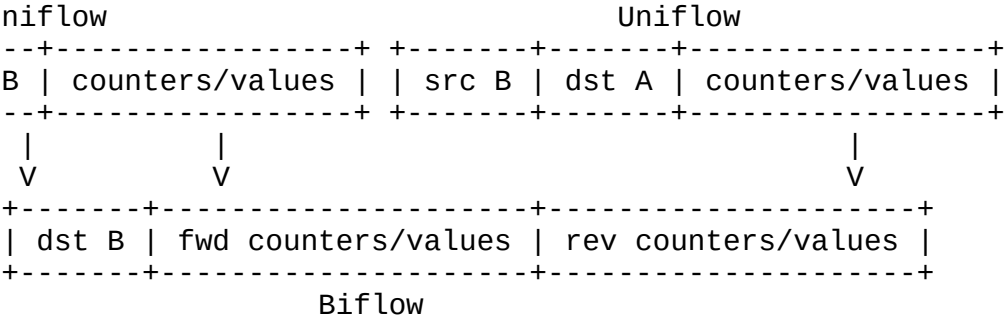

Figure 1: Bidirectional Flow Conceptual Diagram

The reverse direction values are represented by Reverse Information Elements. The representation of these Reverse Information Elements within Templates is detailed in Section 5. A Flow Record may be considered to be a Biflow record by the Collecting Process if it contains at least one Reverse Information Element AND at least one Directional Key Field. Flow Records containing Reverse Information Elements but no Directional Key Fields are illegal, MUST NOT be sent by the Exporting Process, and SHOULD be dropped by the Collecting Process. The Collecting Process SHOULD log the receipt of such illegal Flow Records.

When exporting Uniflows, Exporting Processes SHOULD use a Template containing no Reverse Information Elements. Note that a Template whose only Reverse Information Elements are counters MAY be used to

export Uniflows, as counters with values of 0 are semantically equivalent to no reverse direction. However, this approach is not possible for Reverse Information Elements whose zero values have a distinct meaning (e.g., tcpControlBits).

Note that a Biflow traversing a middlebox [RFC3234] may show different flow properties on each side of the middlebox due to changes to the packet header or payload performed by the middlebox itself. Therefore, it MUST be clear at a Collecting Process whether packets were observed and metered before or after modification. The Observation Process SHOULD be located on one side of a middlebox, and the Exporting Process SHOULD communicate to the Collecting Process both the incoming value of the flow property changed within the middlebox and the changed value on the "other side". The IPFIX Information Model [RFC5102] provides Information Elements with prefix "post" for this purpose. The location of the Observation Point(s) with respect to the middlebox can be communicated using Options with Observation Point as Scope and elements such as lineCardID or samplerID.

For further information on the effect of middleboxes within the IPFIX architecture, refer to Section 7 of the IPFIX Implementation Guidelines [IPFIX-IMPLEMENTATION].

By the definition of Observation Domain in Section 2 of the IPFIX Protocol document [RFC5101], Biflows may be composed only of packets observed within the same Observation Domain. This implies that Metering Processes that build Biflows out of Uniflows must ensure that the two Uniflows were observed within the same Observation Domain.

5. Direction Assignment

Due to the variety of flow measurement applications and restrictions on Metering Process deployment, one single method of assigning the directions of a Biflow will not apply in all cases. This section describes three methods of direction assignment, and recommends them based upon Metering Process position and measurement application requirements. In each of the figures in this section, the "MP" box represents the Metering Process.

As the method selection is dependent on Metering Process position, it is sufficient to configure the direction assignment method at the Collecting and/or the Exporting Process out-of-band. For example, a Collecting Process might be configured that a specific Exporting Process identified by exporterIPv4Address is assigning direction by initiator; or both a Collecting Process and an Exporting Process could be simultaneously configured with a specific direction

assignment perimeter. However, for Exporting Processes that use multiple direction selection methods, or for Collecting Processes accepting data from Exporting Processes using a variety of methods, a biflowDirection Information Element is provided for optional representation of direction assignment information.

5.1. Direction by Initiator

If the measurement application requires the determination of the initiator and responder of a given communication, the Metering Process SHOULD define the Biflow Source to be the initiator of the Biflow, where possible. This can be roughly approximated by a Metering Process observing packets in both directions simply assuming that the first packet seen in a given Biflow is the packet initiating the Biflow. A Metering Process may improve upon this method by using knowledge of the transport or application protocols (e.g., TCP flags, DNS question/answer counts) to better approximate the flow-initiating packet.

Note that direction assignment by initiator is most easily done by a single Metering Process positioned on a local link layer, as in Figure 2, or a single Metering Process observing bidirectional packet flows at a symmetric perimeter routing point, as in Figure 3.

Note also that many Metering Processes have an "active" timeout, such that any flow with a duration longer than the active timeout is expired and any further packets belonging to that flow are accounted for as part of a new flow. This mechanism may cause issues with the assumption that a first packet seen is from the flow initiator, if the "first" packet is a middle packet in a long-duration flow.

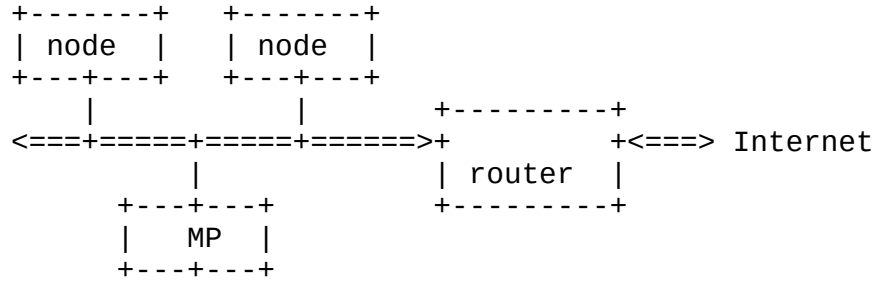

Figure 2: Local Link Metering Process Position

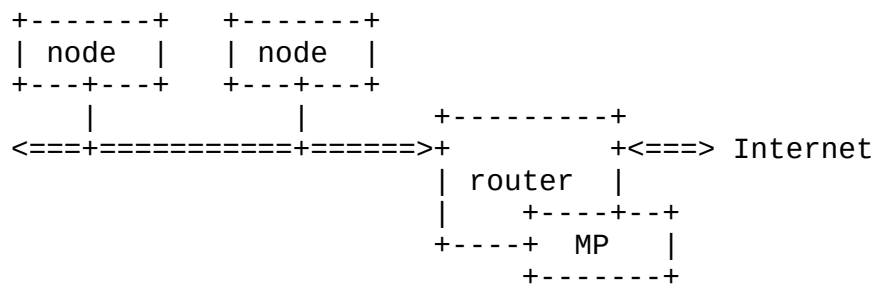

Figure 3: Symmetric Routing Point Metering Process Position

5.2. Direction by Perimeter

If the measurement application is deployed at a network perimeter, as illustrated in Figure 4, such that there is a stable set of addresses that can be defined as "inside" that perimeter, and there is no measurement application requirement to determine the initiator and responder of a given communication, then the Metering Process SHOULD assign the Biflow Source to be the endpoint outside the perimeter.

No facility is provided for exporting the address set defining the interior of a perimeter; this set may be deduced by the Collecting Process observing the set of Biflow Source and Biflow Destination addresses, or configured out-of-band.

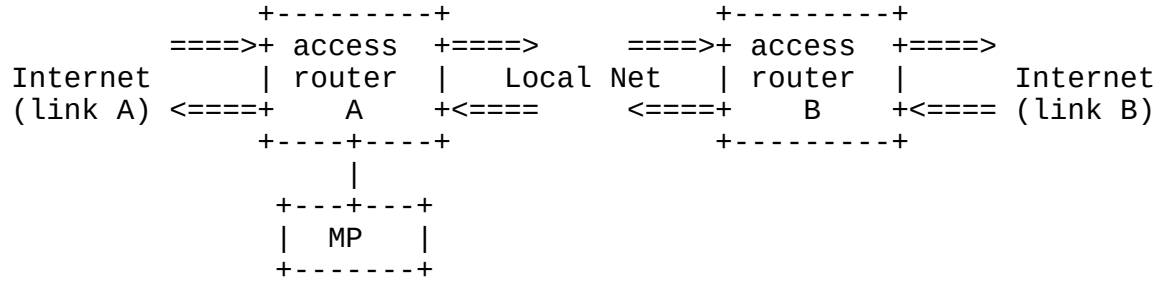

Figure 4: Perimeter Metering Process Position

5.3. Arbitrary Direction

If the measurement application is deployed in a network core, such that there is no stable set of addresses defining a perimeter (e.g., due to BGP updates), as in Figure 5, and no requirement or ability to determine the initiator or responder of a given communication, then the Metering Process MAY assign the Biflow Source and Biflow Destination endpoints arbitrarily.

In this case, the Metering Process SHOULD be consistent in its choice of direction. Once assigned, direction SHOULD be maintained for the lifetime of the Biflow, even in the case of active timeout of a long-lived Biflow.

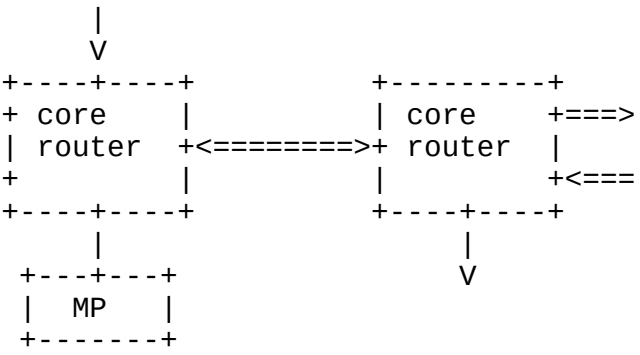

Figure 5: Transit/Core Metering Process Position

6. Record Representation

As noted above, Biflows are exported using a single Flow Record, each of which contains the Flow Key fields once, and both forward Information Elements and Reverse Information Elements for each non-key field. The IPFIX Information Model is extended to provide a Reverse Information Element counterpart to each presently defined forward Information Element, as required by any Information Element that may be a non-key field in a Biflow.

6.1. Reverse Information Element Private Enterprise Number

Reverse Information Elements are specified as a separate "dimension" in the Information Element space, assigning Private Enterprise Number (PEN) 29305 to this document, and defining that PEN to signify "IPFIX Reverse Information Element" (the Reverse PEN). This Reverse PEN serves as a "reverse direction flag" in the Template; each Information Element number within this PEN space is assigned to the reverse counterpart of the corresponding IANA-assigned public Information Element number. In other words, to generate a Reverse Information Element in a Template corresponding to a given forward Information Element, simply set the enterprise bit and define the Information Element within the Reverse PEN space, as in Figure 6 below.

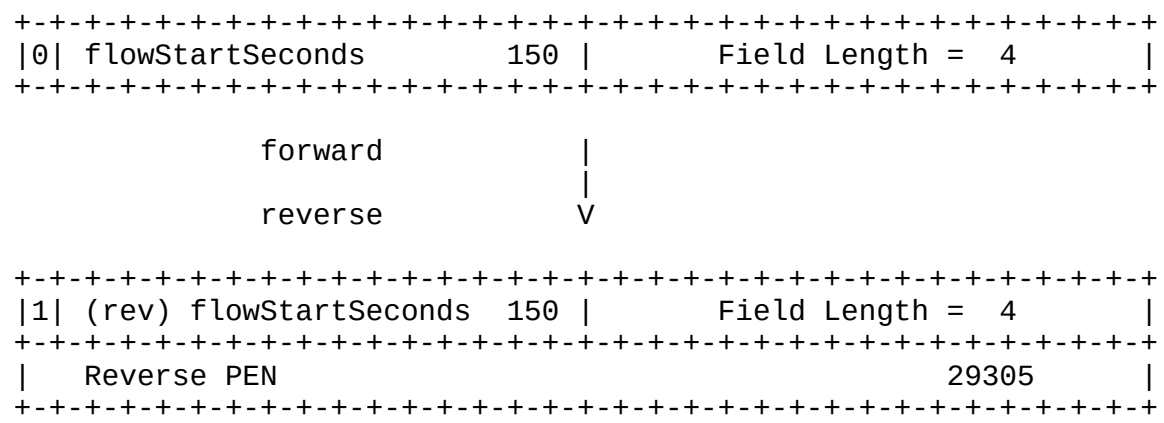

Figure 6: Example Mapping between Forward and Reverse IEs

As the Reverse Information Element dimension is treated explicitly as such, new Information Elements can be added freely to the IANA-managed space without concern for whether a Reverse Information Element should also be added. Aside from the initial allocation of a Private Enterprise Number for this purpose, there is no additional maintenance overhead for supporting Reverse Information Elements in the IPFIX Information Model.

Note that certain Information Elements in the IPFIX Information Model [RFC5102] are not reversible; that is, they are semantically meaningless as Reverse Information Elements. An Exporting Process MUST NOT export a Template containing the reverse counterpart of a non-reversible Information Element. A Collecting Process receiving the reverse counterpart of a non-reversible Information Element MAY discard that Information Element from the Flow Record. Non-reversible Information Elements represent properties of the Biflow record as a whole, or are intended for internal the use of the IPFIX Protocol itself. Therefore, by definition, they cannot be associated with a single direction or endpoint of the Flow.

The following specific Information Elements are not reversible:

1. Identifiers defined in Section 5.1 of [RFC5102] that cannot be associated with a single direction of Uniflow collection: flowId (5.1.7), templateId (5.1.8), observationDomainId (5.1.9), and commonPropertiesId (5.1.11).

2. Process configuration elements defined in Section 5.2 of [RFC5102].

3. Process statistics elements defined in Section 5.3 of [RFC5102].

4. paddingOctets defined in Section 5.12.1 of [RFC5102].

5. biflowDirection (defined in Section 6.3 of this document).

Any future addition to the Information Element Registry by IANA that meets the criteria defined above SHOULD also be considered to be non-reversible by the Collecting Process.

Note that Information Elements commonly used as Flow Keys (e.g., header fields defined in Sections 5.4 and 5.5 of the Information Model) are reversible, as they may be used as value fields in certain contexts, as when associating ICMP error messages with the flows that caused them.

6.2. Enterprise-Specific Reverse Information Elements

Note that the Reverse PEN defined above is only available for allocating reverse counterparts of IANA-registered IPFIX Information Elements. No facility is provided for allocating reverse counterparts of enterprise-specific Information Elements.

The allocation of enterprise-specific Information Elements for IPFIX is left to the discretion of the organization allocating them. Note that, as enterprise-specific Information Elements are designed for the internal use of private enterprises, the lack of any guidance or standard on Information Element allocation policies poses no interoperability issues. However, if a private enterprise's own Information Element registry anticipates the allocation of reversible Information Elements, and the use of this specification for the export of Biflow data, that registry MAY reserve one of the fifteen available bits in the Information Element ID to signify the reverse direction. For example, if the most significant bit were selected, this would reserve Information Element IDs 0x4000 to 0x7FFF for the reverse direction of Information Element IDs 0x0000 to 0x3FFF.

6.3. biflowDirection Information Element

Description: A description of the direction assignment method used to assign the Biflow Source and Destination. This Information Element MAY be present in a Flow Record, or applied to all flows exported from an Exporting Process or Observation Domain using IPFIX Options. If this Information Element is not present in a Flow Record or associated with a Biflow via scope, it is assumed that the configuration of the direction assignment method is done out-of-band. Note that when using IPFIX Options to apply this Information Element to all flows within an Observation Domain or from an Exporting Process, the Option SHOULD be sent reliably. If reliable transport is not available (i.e., when using UDP), this

Information Element SHOULD appear in each Flow Record. This field may take the following values:

Value	Name	Description
0x00	arbitrary	Direction was assigned arbitrarily.
0x01	initiator	The Biflow Source is the flow initiator, as determined by the Metering Process' best effort to detect the initiator.
0x02	reverseInitiator	The Biflow Destination is the flow initiator, as determined by the Metering Process' best effort to detect the initiator. This value is provided for the convenience of Exporting Processes to revise an initiator estimate without re-encoding the Biflow Record.
0x03	perimeter	The Biflow Source is the endpoint outside of a defined perimeter. The perimeter's definition is implicit in the set of Biflow Source and Biflow Destination addresses exported in the Biflow Records.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

ElementId: 239

Status: current

7. IANA Considerations

This document specifies the creation of a new dimension in the Information Element space defined by the IPFIX Information Model [RFC5102]. This new dimension is defined by the allocation of a new Private Enterprise Number (PEN). The Internet Assigned Numbers Authority (IANA) has assigned Private Enterprise Number 29305 to this document as the "IPFIX Reverse Information Element Private Enterprise", with this document's authors as point of contact.

This document specifies the creation of a new IPFIX Information Element, biflowDirection, as defined in Section 6.3. IANA has assigned Information Element number 239 in the IPFIX Information

Element registry for the biflowDirection Information Element. The values defined for this Information Element are static, and as such do not need to be maintained by IANA in a sub-registry.

8. Security Considerations

The same security considerations as for the IPFIX Protocol [RFC5101] apply.

9. Acknowledgments

We would like to thank Lutz Mark, Juergen Quittek, Andrew Johnson, Paul Aitken, Benoit Claise, and Carsten Schmoll for their contributions and comments. Special thanks to Michelle Cotton for her assistance in navigating the IANA process for Enterprise Number assignment, and for the IANA pre-review of the document.

10. References

10.1. Normative References

[RFC5101] Claise, B., Ed., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", RFC 5101, January 2008.

[RFC5102] Quittek, J., Bryant, S., Claise, B., Aitken, P., and J. Meyer, "Information Model for IP Flow Information Export", RFC 5102, January 2008.

10.2. Informative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.

[RFC3917] Quittek, J., Zseby, T., Claise, B., and S. Zander, "Requirements for IP Flow Information Export (IPFIX)", RFC 3917, October 2004.

[IPFIX-ARCH] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", Work in Progress, September 2006.

[IPFIX-AS] Zseby, T., Boschi, E., Brownlee, N., and B. Claise, "IPFIX Applicability", Work in Progress, July 2007.

[IPFIX-IMPLEMENTATION] Boschi, E., Mark, L., Quittek, j., Stiemerling, M., and P. Aitken, "IPFIX Implementation Guidelines", Work in Progress, September 2007.

[IPFIX-REDUCING] Boschi, E., Mark, L., and B. Claise, "Reducing Redundancy in IP Flow Information Export (IPFIX) and Packet Sampling (PSAMP) Reports", Work in Progress, May 2007.

Appendix A. Examples

The following example describes a Biflow record as specified in Section 6, above. The Reverse PEN is assigned for the purpose of differentiating forward from Reverse Information Elements.

The information exported in this case is:

- The start time of the flow: flowStartSeconds in the IPFIX Information Model [RFC5102], with a length of 4 octets.

- The reverse start time of the flow: flowStartSeconds in the IPFIX Information Model [RFC5102], with a length of 4 octets, and the enterprise bit set to 1. The following PEN is the Reverse PEN.

- The IPv4 source IP address: sourceIPv4Address in the IPFIX Information Model [RFC5102], with a length of 4 octets.

- The IPv4 destination IP address: destinationIPv4Address in the IPFIX Information Model [RFC5102], with a length of 4 octets.

- The source port: sourceTransportPort in the IPFIX Information Model [RFC5102], with a length of 2 octets.

- The destination port: destinationTransportPort in the IPFIX Information Model [RFC5102], with a length of 2 octets.

- The protocol identifier: protocolIdentifier in the IPFIX Information Model [RFC5102], with a length of 1 octet.

- The number of octets of the Flow: octetTotalCount in the IPFIX Information Model [RFC5102], with a length of 4 octets.

- The reverse number of octets of the Flow: octetTotalCount in the IPFIX Information Model [RFC5102], with a length of 4 octets, and the enterprise bit set to 1. The following PEN is the Reverse PEN.

- The number of packets of the Flow: packetTotalCount in the IPFIX Information Model [RFC5102], with a length of 4 octets.

- The reverse number of packets of the Flow: packetTotalCount in the IPFIX Information Model [RFC5102], with a length of 4 octets, and the enterprise bit set to 1. The following PEN is the Reverse PEN.

and the resulting Template Set would look like the diagram below:

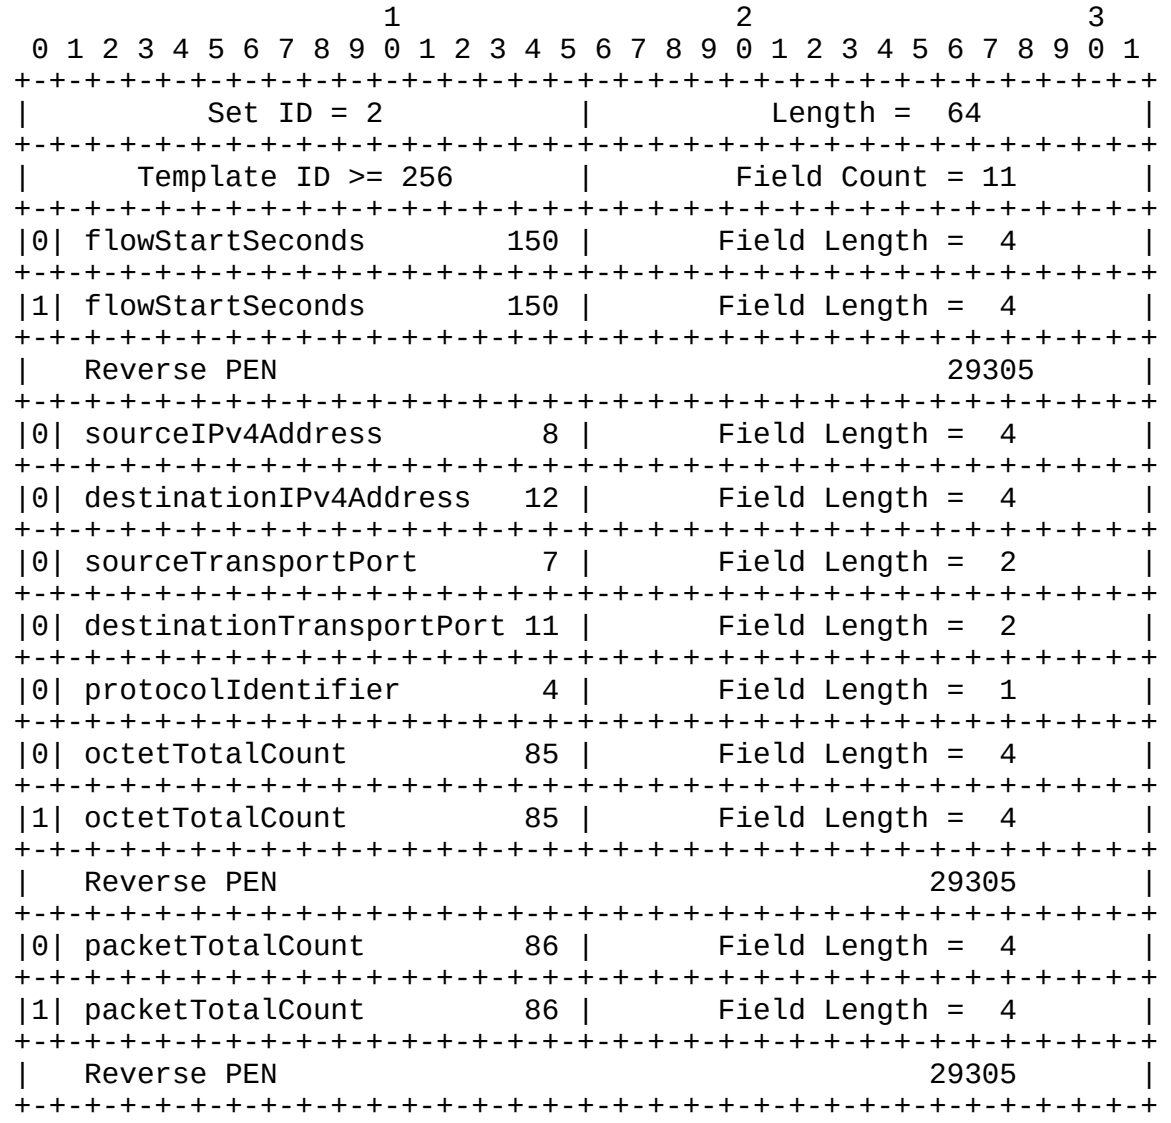

Figure 7: Single Record Biflow Template Set

The following example Data Set represents a typical HTTP transaction. Its format is defined by the example Template, above.

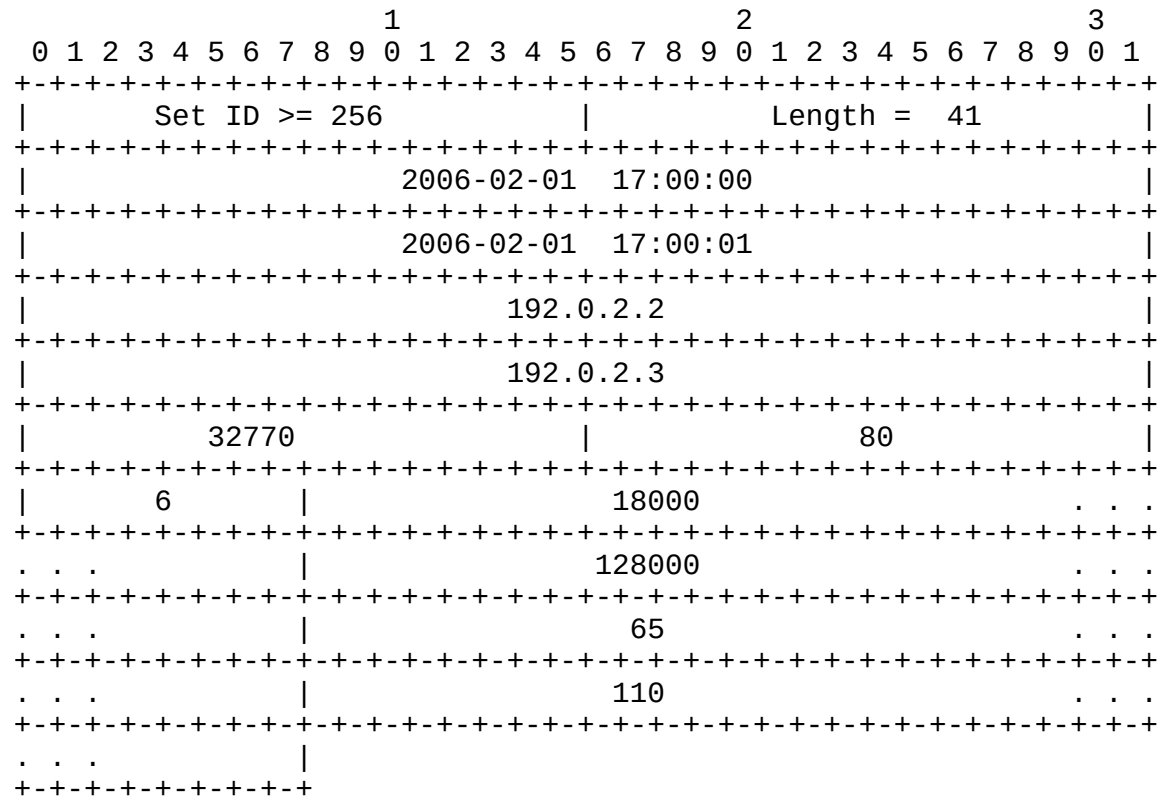

Figure 8: Single Record Biflow Data Set

The following example demonstrates the use of the biflowDirection Information Element, as specified in Section 6.2, using the IPFIX Options mechanism to specify that perimeter direction selection is in effect for a given Observation Domain.

The information exported in this case is:

- The Observation Domain: observationDomainId in the IPFIX Information Model [RFC5102], with a length of 4 octets.

- The direction assignment method: biflowDirection as defined in Section 6.2, above, with a length of 1 octet.

and the resulting Options Template Set would look like the diagram below:

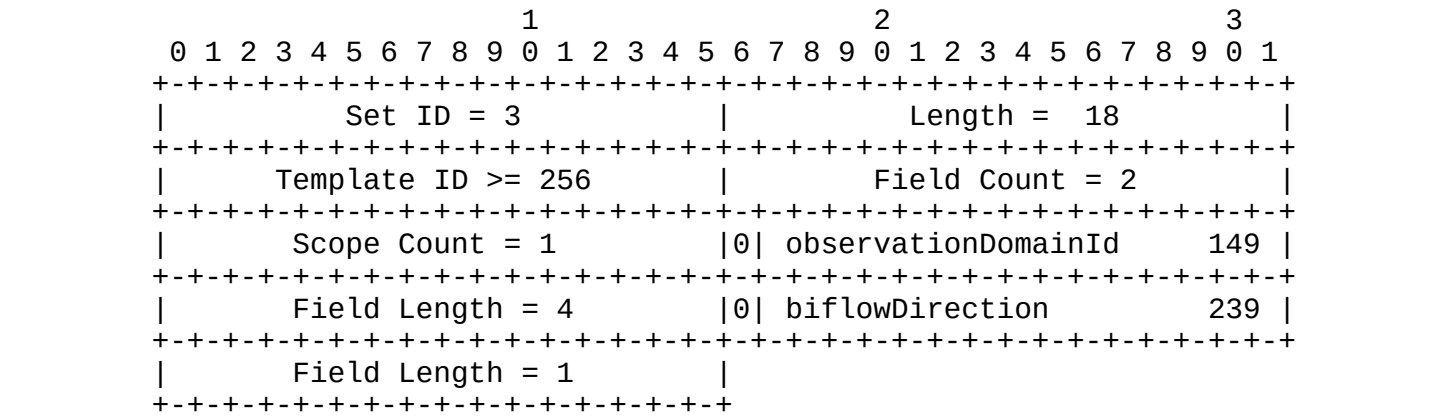

Figure 9: Biflow Direction Options Template Set

The following example Data Set would specify that perimeter direction selection is in effect for the Observation Domain with ID 33. Its format is defined by the example Options Template, above. Note that this example data set would be sent reliably, as specified in the description of the biflowDirection Information Element.

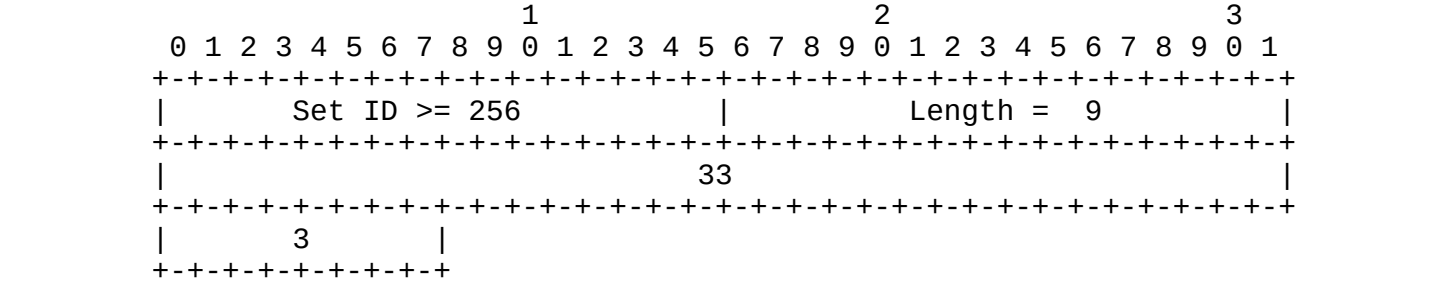

Figure 10: Biflow Direction Options Data Set

Appendix B. XML Specification of biflowDirection Information Element

This appendix contains a machine-readable description of the biflowDirection information element defined in this document, coded in XML. Note that this appendix is of informational nature, while the text in Section 6.3 is normative.

The format in which this specification is given is described by the XML Schema in Appendix B of the IPFIX Information Model [RFC5102].

```
<?xml version="1.0" encoding="UTF-8"?>

<fieldDefinitions xmlns="urn:ietf:params:xml:ns:ipfix-info"
             xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
             xsi:schemaLocation="urn:ietf:params:xml:ns:ipfix-info
             ipfix-info.xsd">

  <field name="biflowDirection" dataType="unsigned8"
         dataTypeSemantics="identifier" group="misc"
         elementId="239" applicability="all" status="current">
    <description>
      <paragraph>
       A description of the direction assignment method used to
       assign the Biflow Source and Destination.  This
       Information Element MAY be present in a Flow Data Record, or
       applied to all flows exported from an Exporting Process or
       Observation Domain using IPFIX Options.  If this Information
       Element is not present in a Flow Record or associated with a
       Biflow via scope, it is assumed that the configuration of
       the direction assignment method is done out-of-band.  Note
       that when using IPFIX Options to apply this Information
       Element to all flows within an Observation Domain or from an
       Exporting Process, the Option SHOULD be sent reliably.  If
       reliable transport is not available (i.e., when using UDP),
       this Information Element SHOULD appear in each Flow
       Record.  This field may take the following values:
           </paragraph>
```

```
            <artwork>
   +-------+------------------+----------------------------------------+
   | Value | Name             | Description                            |
   +-------+------------------+----------------------------------------+
   | 0x00  | arbitrary        | Direction was assigned arbitrarily.    |
   | 0x01  | initiator        | The Biflow Source is the flow          |
   |       |                  | initiator, as determined by the        |
   |       |                  | Metering Process' best effort to       |
   |       |                  | detect the initiator.                  |
   | 0x02  | reverseInitiator | The Biflow Destination is the flow     |
   |       |                  | initiator, as determined by the        |
   |       |                  | Metering Process' best effort to       |
   |       |                  | detect the initiator.  This value is   |
   |       |                  | provided for the convenience of        |
   |       |                  | Exporting Processes to revise an       |
   |       |                  | initiator estimate without re-encoding |
   |       |                  | the Biflow Record.                     |
   | 0x03  | perimeter        | The Biflow Source is the endpoint      |
   |       |                  | outside of a defined perimeter.  The   |
   |       |                  | perimeter's definition is implicit in  |
   |       |                  | the set of Biflow Source and Biflow    |
   |       |                  | Destination addresses exported in the  |
   |       |                  | Biflow Records.                        |
   +-------+------------------+----------------------------------------+
            </artwork>
     </description>
   </field>
</fieldDefinitions>
```

Authors' Addresses

Brian H. Trammell
CERT Network Situational Awareness
Software Engineering Institute
4500 Fifth Avenue
Pittsburgh, PA 15213
United States

Phone: +1 412 268 9748
EMail: bht@cert.org

Elisa Boschi
Hitachi Europe
c/o ETH Zurich
Gloriastrasse 35
8092 Zurich
Switzerland

Phone: +41 44 6327057
EMail: elisa.boschi@hitachi-eu.com

Full Copyright Statement

Intellectual Property