PSAMP working group
Internet Draft
draft-ietf-psamp-protocol-07.txt
Expires: April 2006

EDITOR: B. Claise
Cisco Systems
October 2006

Packet Sampling (PSAMP) Protocol Specifications

Status of this Memo

Copyright Notice

Abstract

This document specifies the export of packet information from a PSAMP Exporting Process to a PSAMP Collecting Process. For export of packet information the IP Flow Information eXport (IPFIX) protocol is used, as both the IPFIX and PSAMP architecture match very well and the means provided by the IPFIX protocol are

sufficient. The document specifies in detail how the IPFIX protocol is used for PSAMP export of packet information.

Conventions used in this document

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].

Table of Contents

1. Introduction

The name PSAMP is a contraction of the phrase Packet SAMPling. The word "sampling" captures the idea that only a subset of all packets passing a network element will be selected for reporting. PSAMP selection operations include random selection, deterministic selection, and deterministic approximations to random selection (hash-based selection).

The IP Flow information export (IPFIX) protocol specified in [IPFIX-PROTO] exports IP traffic information [IPFIX-INFO] observed at network devices. This matches the general protocol requirements outlined in the PSAMP framework [PSAMP-FMWK]. However, there are some architectural differences between IPFIX and PSAMP in the requirements for an export protocol. While the IPFIX architecture [IPFIX-ARCH] is focused on gathering and exporting IP traffic flow information, the focus of the PSAMP framework [PSAMP-FMWK] is on exporting information on individual packets. This basic difference and a set of derived differences in protocol requirements are outlined in Section 4. Despite these differences, the IPFIX protocol is well suited as PSAMP protocol. Section 5 specifies how the IPFIX protocol is used for the export of packet samples. Required extensions of the IPFIX information model are specified in the PSAMP information model [PSAMP-INFO].

2. PSAMP Documents Overview

[PSAMP-FMWK]: "A Framework for Packet Selection and Reporting", describes the PSAMP framework for network elements to select subsets of packets by statistical and other methods, and to export a stream of reports on the selected packets to a collector.

[PSAMP-TECH]: "Sampling and Filtering Techniques for IP Packet Selection", describes the set of packet selection techniques supported by PSAMP.

This document: "Packet Sampling (PSAMP) Protocol Specifications" specifies the export of packet information from a PSAMP Exporting Process to a PSAMP Collecting Process.

[PSAMP-INFO]: "Information Model for Packet Sampling Exports" defines an information and data model for PSAMP.

3. Terminology

As the IPFIX export protocol is used to export the PSAMP information, the relevant IPFIX terminology from [IPFIX-PROTO] is copied over in this document. The terminology summary table in section 4.1 gives a quick overview of the relationships between the different IPFIX terms. The PSAMP terminology defined here is fully consistent with all terms listed in [PSAMP-TECH] and [PSAMP-FMWK] but only definitions that are relevant to the PSAMP protocol appear here. Section 5.4 applies the PSAMP terminology to the IPFIX protocol terminology.

3.1 IPFIX Terminology

The IPFIX terminology section has been entirely copied over from [IPFIX-PROTO], except for the IPFIX Exporting Process term, which is defined more precisely in the PSAMP terminology section.

Observation Point

An Observation Point is a location in the network where IP packets can be observed. Examples include: a line to which a probe is attached, a shared medium, such as an Ethernet-based LAN, a single port of a router, or a set of interfaces (physical or logical) of a router.

Note that every Observation Point is associated with an Observation Domain (defined below), and that one Observation Point may be a superset of several other Observation Points. For example one Observation Point can be an entire line card. That would be the superset of the individual Observation Points at the line card's interfaces.

Observation Domain

An Observation Domain is the largest set of Observation Points for which Flow information can be aggregated by a Metering Process. Each Observation Domain presents itself using a unique ID to the Collecting Process to identify the IPFIX Messages it generates. For

example, a router line card may be an observation domain if it is composed of several interfaces, each of which is an Observation Point. Every Observation Point is associated with an Observation Domain.

IP Traffic Flow or Flow

There are several definitions of the term 'flow' being used by the Internet community. Within the context of IPFIX we use the following definition:

A Flow is defined as a set of IP packets passing an Observation Point in the network during a certain time interval. All packets belonging to a particular Flow have a set of common properties. Each property is defined as the result of applying a function to the values of:

1. one or more packet header field (e.g. destination IP address), transport header field (e.g. destination port number), or application header field (e.g. RTP header fields [RFC1889])

2. one or more characteristics of the packet itself (e.g. number of MPLS labels, etc...)

3. one or more of fields derived from packet treatment (e.g. next hop IP address, the output interface, etc...)

A packet is defined to belong to a Flow if it completely satisfies all the defined properties of the Flow.

This definition covers the range from a Flow containing all packets observed at a network interface to a Flow consisting of just a single packet between two applications. It includes packets selected by a sampling mechanism.

Flow Key

Each of the fields which
1. Belong to the packet header (e.g. destination IP address)
2. Are a property of the packet itself (e.g. packet length)
3. Are derived from packet treatment (e.g. AS number)
and which are used to define a Flow are termed Flow Keys.

Flow Record

A Flow Record contains information about a specific Flow that was observed at an Observation Point. A Flow Record contains measured properties of the Flow (e.g. the total number of bytes for all the Flow's packets) and usually characteristic properties of the Flow (e.g. source IP address).

Metering Process

The Metering Process generates Flow Records. Inputs to the process are packet headers and characteristics observed at an Observation Point, and packet treatment at the Observation Point (for example the selected output interface).

The Metering Process consists of a set of functions that includes packet header capturing, timestamping, sampling, classifying, and maintaining Flow Records.

The maintenance of Flow Records may include creating new records, updating existing ones, computing Flow statistics, deriving further Flow properties, detecting Flow expiration, passing Flow Records to the Exporting Process, and deleting Flow Records.

Exporter

A device which hosts one or more Exporting Processes is termed an Exporter.

IPFIX Device

An IPFIX Device hosts at least one Observation Point, a Metering Process and an Exporting Process.

Collecting Process

A Collecting Process receives Flow Records from one or more Exporting Processes. The Collecting Process might process or store received Flow Records, but such actions are out of scope for this document.

Collector

A device which hosts one or more Collecting Processes is termed a Collector.

Template

A Template is an ordered sequence of <type, length> pairs, used to completely specify the structure and semantics of a particular set of information that needs to be communicated from an IPFIX Device to a Collector. Each Template is uniquely identifiable by means of a Template ID.

IPFIX Message

An IPFIX Message is a message originating at the Exporting Process that carries the IPFIX records of this Exporting Process and whose destination is a Collecting Process. An IPFIX Message is encapsulated at the transport layer.

Message Header

The Message Header is the first part of an IPFIX Message, which provides basic information about the message such as the IPFIX version, length of the message, message sequence number, etc.

Template Record

A Template Record defines the structure and interpretation of fields in a Data Record.

Data Record

A Data Record is a record that contains values of the parameters corresponding to a Template Record.

Options Template Record

An Options Template Record is a Template Record that defines the structure and interpretation of fields in a Data Record, including defining how to scope the applicability of the Data Record.

Set

Set is a generic term for a collection of records that have a similar structure. In an IPFIX Message, one or more Sets follow the Message Header.

There are three different types of Sets: Template Set, Options Template Set, and Data Set.

Template Set

A Template Set is a collection of one or more Template Records that have been grouped together in an IPFIX Message.

Options Template Set

An Options Template Set is a collection of one or more Options Template Records that have been grouped together in an IPFIX Message.

Data Set

A Data Set is one or more Data Records, of the same type, that are grouped together in an IPFIX Message. Each Data Record is previously defined by a Template Record or an Options Template Record.

Information Element

An Information Element is a protocol and encoding independent description of an attribute which may appear in an IPFIX Record. The IPFIX information model [IPFIX-INFO] defines the base set of Information Elements for IPFIX. The type associated with an Information Element indicates constraints on what it may contain and also determines the valid encoding mechanisms for use in IPFIX.

	contents	
Set	Template	record
Data Set	/	Data Record(s)
Template Set	Template Record(s)	/
Options Template Set	Options Template Record(s)	/

Figure A: Terminology Summary Table

3.2 PSAMP Terminology

The PSAMP terminology section has been copied over from [PSAMP-TECH].

3.2.1 Packet Streams and Packet Content

Observed Packet Stream

The Observed Packet Stream is the set of all packets observed at the Observation Point.

Packet Stream

A packet stream denotes a subset of the Observed Packet Stream that flows past some specified point within the Selection Process. An example of a Packet Stream is the output of the Selection Process. Note that packets selected from a stream, e.g. by Sampling, do not necessarily possess a property by which they can be distinguished

from packets that have not been selected. For this reason the term "stream" is favored over "flow", which is defined as set of packets with common properties [RFC3917].

Packet Content

The packet content denotes the union of the packet header (which includes link layer, network layer and other encapsulation headers) and the packet payload.

3.2.2 Selection Process

Selection Process

A Selection Process takes the Observed Packet Stream as its input and selects a subset of that stream as its output.

Selection State

A Selection Process may maintain state information for use by the Selection Process. At a given time, the Selection State may depend on packets observed at and before that time, and other variables. Examples include:

(i) sequence numbers of packets at the input of Selectors;

(ii) a timestamp of observation of the packet at the Observation Point;

(iii) iterators for pseudorandom number generators;

(iv) hash values calculated during selection;

(v) indicators of whether the packet was selected by a given Selector.

Selection Processes may change portions of the Selection State as a result of processing a packet. Selection state for a packet is to reflect the state after processing the packet.

Selector

A Selector defines the action of a Selection Process on a single packet of its input. If selected, the packet becomes an element of the output Packet Stream.

The Selector can make use of the following information in determining whether a packet is selected:

(i) the Packet Content;

(ii) information derived from the packet's treatment at the Observation Point;

(iii) any selection state that may be maintained by the Selection Process.

Composite Selector

A Composite Selector is an ordered composition of Selectors, in which the output Packet Stream issuing from one Selector forms the input Packet Stream to the succeeding Selector.

Primitive Selector

A Selector is primitive if it is not a Composite Selector.

Selector ID

The Selector ID is the unique ID identifying a Primitive Selector. The ID is unique within the Observation Domain.

Selection Sequence

From all the packets observed at an Observation Point, only a few packets are selected by one or more Selectors. The Selection Sequence is a unique value per Observation Domain describing the Observation Point and the Selector IDs through which the packets are selected.

3.2.3 Reporting

Packet Reports

Packet Reports comprise a configurable subset of a packet's input to the Selection Process, including the Packet Content, information relating to its treatment (for example, the output interface), and its associated selection state (for example, a hash of the Packet Content)

Report Interpretation

Report Interpretation comprises subsidiary information, relating to one or more packets, that are used for interpretation of their Packet Reports. Examples include configuration parameters of the Selection Process.

Report Stream

The Report Stream is the output of a Selection Process, comprising two distinguished types of information: Packet Reports, and Report Interpretation.

3.2.4 Exporting Process

Exporting Process

An Exporting Process sends, in the form of Export Packets, the output of one or more Metering Processes to one or more Collectors.

Export Packet

An Export Packet is a combination of Report Interpretation(s) and/or one or more Packet Reports that are bundled by the Exporting Process into a Export Packet for exporting to a Collector.

3.2.5 PSAMP Device

PSAMP Device

A PSAMP Device is a device hosting at least an Observation Point, a Selection Process and an Exporting Process. Typically, corresponding Observation Point(s), Selection Process(es) and Exporting Process(es) are co-located at this device, for example at a router.

3.2.6 Selection Methods

Filtering

A filter is a Selector that selects a packet deterministically based on the Packet Content, or its treatment, or functions of these occurring in the Selection State. Examples include field match Filtering, and Hash-based Selection.

Sampling

A Selector that is not a filter is called a Sampling operation. This reflects the intuitive notion that if the selection of a packet

cannot be determined from its content alone, there must be some type of Sampling taking place.

Content-independent Sampling

A Sampling operation that does not use Packet Content (or quantities derived from it) as the basis for selection is called a Content-independent Sampling operation. Examples include systematic Sampling, and uniform pseudorandom Sampling driven by a pseudorandom number whose generation is independent of Packet Content. Note that in Content-independent Sampling it is not necessary to access the Packet Content in order to make the selection decision.

Content-dependent Sampling

A Sampling operation where selection is dependent on Packet Content is called a Content-dependent Sampling operation. Examples include pseudorandom selection according to a probability that depends on the contents of a packet field. Note that this is not a filter, because the selection is not deterministic.

Hash Domain

A subset of the Packet Content and the packet treatment, viewed as an N-bit string for some positive integer N.

Hash Range

A set of M-bit strings for some positive integer M that define the range of values the result of the hash operation can take.

Hash Function

A deterministic map from the Hash Domain into the Hash Range.

Hash Selection Range

A subset of the Hash Range. The packet is selected if the action of the Hash Function on the Hash Domain for the packet yields a result in the Hash Selection Range.

Hash-based Selection

Filtering specified by a Hash Domain, a Hash Function, a Hash Range and a Hash Selection Range.

Approximative Selection

Selectors in any of the above categories may be approximated by operations in the same or another category for the purposes of implementation. For example, uniform pseudorandom Sampling may be approximated by Hash-based Selection, using a suitable Hash Function and Hash Domain. In this case, the closeness of the approximation depends on the choice of Hash Function and Hash Domain.

Population

A Population is a Packet Stream, or a subset of a Packet Stream. A Population can be considered as a base set from which packets are selected. An example is all packets in the Observed Packet Stream that are observed within some specified time interval.

Population Size

The Population Size is the number of all packets in the Population.

Sample Size

The number of packets selected from the Population by a Selector.

Configured Selection Fraction

The Configured Selection Fraction is the ratio of the number of packets selected by a Selector from an input Population, to the Population Size, as based on the configured selection parameters.

Attained Selection Fraction

The Attained Selection Fraction is the actual ratio of the number of packets selected by a Selector from an input Population, to the Population Size. For some Sampling methods the Attained Selection Fraction can differ from the Configured Selection Fraction due to, for example, the inherent statistical variability in Sampling decisions of probabilistic Sampling and Hash-based Selection. Nevertheless, for large Population Sizes and properly configured Selectors, the Attained Selection Fraction usually approaches the Configured Selection Fraction.

3.3 IPFIX and PSAMP Terminology Comparison

The PSAMP terminology has been specified with an IPFIX background, as PSAMP and IPFIX have similar terms. However, this section clarifies the terms between the IPFIX and PSAMP terminology.

3.3.1 IPFIX and PSAMP Processes

The figure B indicates the sequence of the processes (Metering and Exporting) within the PSAMP Device.

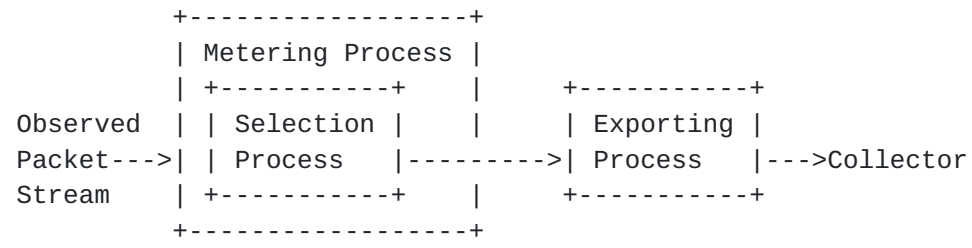

Figure B: PSAMP Processes

The Selection Process, which takes an Observed Packet Stream as its input, is an integral part of the Metering Process. The Selection Process chooses which packets from its input packet stream will be reported on by the rest of the Metering Process. Note that a "Process" is not necessarily implemented as a separate CPU thread.

3.3.2 Packet Report, Packet Interpretation, and Data Record

The PSAMP terminology speaks of Packet Report and Packet Interpretation, while the IPFIX terminology speaks of Data Record and (Option) Template Record. The PSAMP Packet Report, which comprises information about the observed packet, can be viewed as analogous to the IPFIX Data Record defined by a Template Record. The PSAMP Packet Interpretation, which comprises subsidiary information used for the interpretation of the Packet Reports, can be viewed as analogous to the IPFIX Data Record defined by an Option Template Record.

4. Differences between PSAMP and IPFIX

The output of the IPFIX working group relevant for this draft is structured into three documents:

- IP Flow information architecture [IPFIX-ARCH]
- IPFIX protocol specifications [IPFIX-PROTO]
- IP Flow information export information model [IPFIX-INFO]

In the following sections we investigate the differences between IPFIX and PSAMP for each of those aspects.

4.1 Architecture Point of View

Traffic Flow measurement as described in the IPFIX requirements [RFC3917] and the IPFIX architecture [IPFIX-ARCH] can be separated into two stages: packet processing and Flow processing. Figure C illustrates these stages.

In stage 1, all processing steps act on packets. Packets are captured, time stamped, selected by one or more selection steps and finally forwarded to packet classification that maps packets to Flows. The packets selection steps may include Filtering and Sampling functions.

In stage 2, all processing steps act on Flows. After packets are classified (mapped to Flows), Flows are generated (or updated if they exist already). Flow generation and update steps may be performed repeatedly for aggregating Flows. Finally, Flows are exported.

Packet Sampling as described in the PSAMP framework [PSAMP-FMWK] covers only stage 1 of the IPFIX architecture with the packet classification replaced by packet record export.

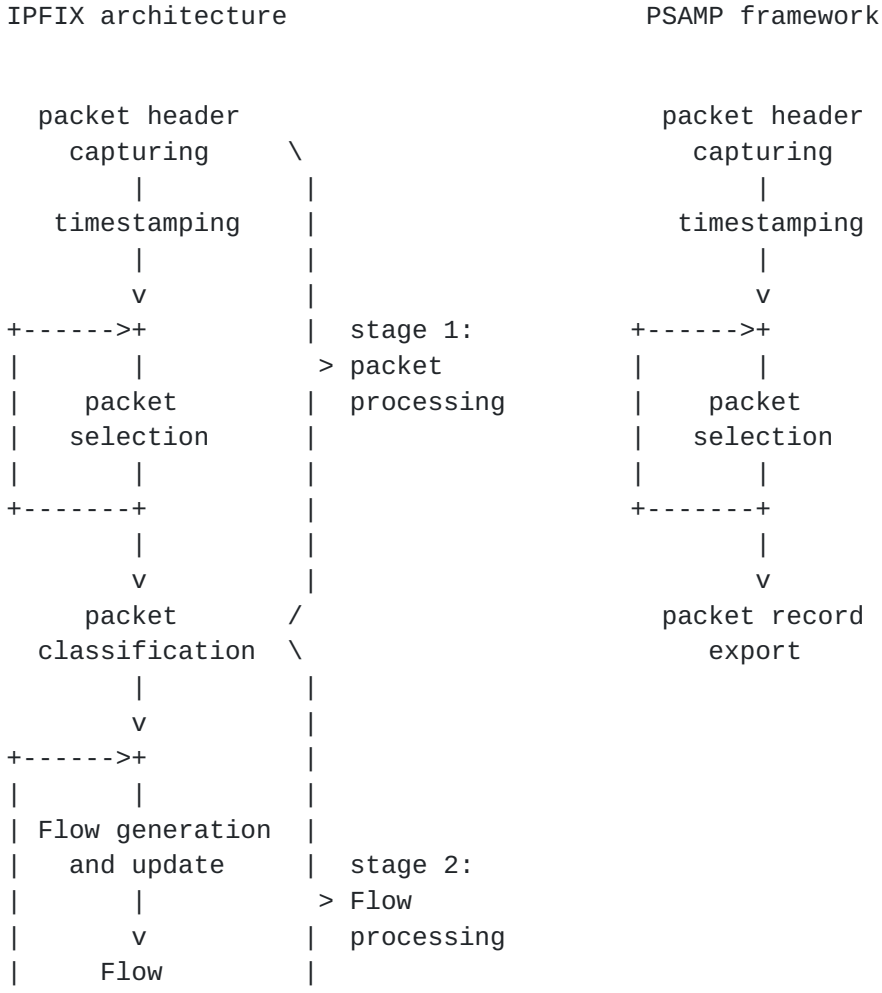

```
     |   selection      |
```

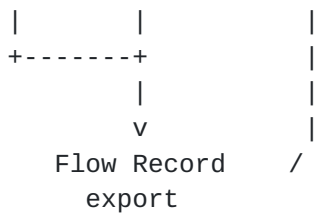

Figure C: Comparison of IPFIX architecture and PSAMP framework

4.2 Protocol Point of View

Concerning the protocol, the major difference between IPFIX and PSAMP is that the IPFIX protocol exports Flow Records while the PSAMP protocol exports Packet Records. From a pure export point of view, IPFIX will not distinguish a Flow Record composed of several packets aggregated together, from a Flow Record composed of a single packet. So the PSAMP export can be seen as special IPFIX Flow Record containing information about a single packet.

All extensions of the IPFIX protocol that are required to satisfy the PSAMP requirements have already been incorporated in the IPFIX protocol [IPFIX-PROTO], which was developed in parallel with the PSAMP protocol. An example is the need for a data type for protocol fields that have flexible length, such as an octet array. This was added to the IPFIX protocol specification in order to meet the requirement of the PSAMP protocol to report content of captured packets, for example the first octets of a packet.

4.3 Information Model Point of View

From the information model point of view, the overlap between both the IPFIX and PSAMP protocols is quite large. Most of the Information Elements in the IPFIX protocol are also relevant for exporting packet information, for example all fields reporting packet header properties. Only a few Information Elements, such as observedFlowTotalCount (whose value will always be 1 for PSAMP) etc., cannot be used in a meaningful way by the PSAMP protocol. Also, IPFIX protocol requirements concerning stage 2 of figure C do not apply to the PSAMP metering process.

Further required extensions apply to the information model. Even if the IPFIX charter speaks of Sampling, no Sampling related Information Elements are specified in [IPFIX-INFO]. The task of specifying them was intentionally left for the PSAMP information model [PSAMP-INFO]. A set of several additional fields is required for satisfying the requirements for the PSAMP information model [PSAMP-TECH].

Exploiting the extensibility of the IPFIX information model, the required extension is covered by the PSAMP information model specified in [PSAMP-INFO].

5. PSAMP Requirements versus the IPFIX Solution

In the "Generic Requirements for PSAMP" section, [PSAMP-FMWK] describes some requirements that affect directly the PSAMP export protocol.

In the "Generic Selection Process Requirements" section, [PSAMP-FMWK] describes one requirement that, if not directly related to the export protocol, will put some constraints on it. Parallel Measurements: multiple independent selection processes at the same entity.

Finally, [PSAMP-FMWK] describes a series of requirements specifying the different Information Elements that MUST and SHOULD be reported to the Collector. Nevertheless IPFIX, being a generic export protocol, can export any Information Elements as long as they are described in the information model. So these requirements are mainly targeted for the [PSAMP-INFO] document.

The PSAMP protocol specifications meets almost all the protocol requirements stated in the PSAMP framework document [PSAMP-FMWK]:

* Extensibility
* Parallel selection processes
* Encrypted packets
* Indication of information loss
* Accuracy
* Privacy
* Timeliness
* Congestion avoidance
* Secure export
* Export rate limit
* Microsecond timestamp resolution

The only requirement that is not met is Export Packet compression. With the choice of IPFIX as PSAMP export protocol, the export packet compression option mentioned in the section 8.5 of the framework document [PSAMP-FMWK] is not addressed.

5.1 High Level View of the Integration

The Template Record in the Template Set is used to describe the different PSAMP Information Elements that will be exported to the Collector. The Collector decodes the Template Record in the Template Set and knows which Information Elements to expect when it receives

the Data Records in the PSAMP Packet Report Data Set. Typically, in the base level of the PSAMP functionality, the Template Set will contain the input sequence number, the packet fragment (some number of contiguous bytes from the start of the packet or from the start of the payload) and the Selection Sequence.

The Options Template Record in the Options Template Set is used to describe the different PSAMP Information Elements that concern the Metering Process itself: Sampling and/or Filtering functions, and the associated parameters. The Collector decodes the Options Template Records in the Option Template Set and knows which Information Elements to expect when it receives the Data Records in the PSAMP Report Interpretation Data Set. Typically, the Options Template would contain the Selection Sequence, the Sampling or Filtering functions, and the Sampling or Filtering associated parameters.

PSAMP requires all the different possibilities of the IPFIX protocol specifications [IPFIX-PROTO]. That is the 3 types of Set (Data Set, Template Set and Options Templates Set) with the 2 types of Templates Records (Template Record and Options Template Record), as described in the figure A. As a consequence, PSAMP can't rely on a subset of the IPFIX protocol specifications described in [IPFIX-PROTO]. The entire IPFIX protocol specifications [IPFIX-PROTO] MUST be implemented for the PSAMP protocol.

6. Using the IPFIX Protocol for PSAMP

In this section, we describe the usage of the IPFIX protocol for PSAMP. We describe the record formats and the additional requirements that must be met. PSAMP uses two different types of messages:

- Packet Reports
- Report Interpretation

The format of Packet Reports is defined in IPFIX Template Records. The PSAMP data is transferred as Information Elements in IPFIX Data Records as described by the Template Record. There are two different types of Packet Reports. Basic Packet Reports contain only the basic Information Elements required for PSAMP reporting. Extended Packet Reports MAY contain further Information Elements.
The format of Report Interpretations is defined in IPFIX Option Template Record. The Information Elements are transferred in IPFIX Data Records as described by the Option Template Record. There are four different types of Report Interpretation messages:

- Selection Sequence Report Interpretation
- Selector Report Interpretation
- Selection Sequence Statistics Report Interpretation
- Accuracy Report Interpretation

A description and examples about the usage of those reports is given below.

6.1 Selector ID

The Selector ID is the unique ID identifying a Primitive Selector. Each Primitive Selector MUST have a unique ID within the Observation Domain. The Selector ID is represented by the selectorId Information Element [PSAMP-INFO].

6.2 The Selection Sequence ID

From all the packets observed at an Observation Point, a subset of packets is selected by one or more Selectors. The Selection Sequence is the combination of an Observation Point and one or more Selector(s) through which the packets are selected. The Selection Sequence ID is a unique value representing that combination. The Selection Sequence ID is represented by the selectionSequenceId Information Element [PSAMP-INFO].

6.3 The Exporting Process

An Exporting Process MUST be able to limit the export rate according to a configurable value. The Exporting Process MAY limit the export rate on a per Collecting Process basis.

6.4 Packet Report

For each Selection Sequence, for each selected packet, a Packet Report MUST be created. The format of the Packet Report is specified in a Template Record contained in a Template Set.

There are two types of Packet Report, as described in [PSAMP-FMWK]: the basic Packet Report and the extended Packet Report.

6.4.1 Basic Packet Report

For each selected packet, the Packet Report MUST contain the following information:
- The selectionSequenceId Information Element
If there is a digest function in the selection sequence, the Packet report MUST contain the hash value (digestHashValue Information Element) generated by the digest hash function for each selected packet. If there is more than one digest function then each hash value MUST be included in the same order as they appear in the selection sequence. If there are no digest functions in the selection sequence no element for the digest needs to be sent.

- Some number of contiguous bytes from the start of the packet, including the packet header (which includes link layer, network layer and other encapsulation headers) and some subsequent bytes of the packet payload. Alternatively, the number of contiguous bytes may start at the beginning of the payload. The dataLinkFrameSection, mplsLabelStackSection, mplsPayloadPacketSection, ipPacketSection, and ipPayloadPacketSection PSAMP Information Elements are available for this use. If one of those Information Elements that contain some number of contiguous bytes has got a content with an insufficient number of octets compared to its length specified in the Template, then this Information Element MUST be sent with a new Template using either a fixed length Information Element of the necessary size or a variable length Information Element.

For each selected packet, the Packet Report SHOULD contain the following information:
- the observationTimeMicroseconds Information Element

In the Packet Report, the PSAMP device MUST be capable of exporting the number of observed packets and the number of packets selected by each instance of its Primitive Selectors (as described by the non scope Information Elements of the Selection Sequence Statistics Report Interpretation) although it MAY be a configurable option not to include them. If exported, the Attained Selection Fraction may be calculated precisely for the Observed Packet Stream. The Packet Report MAY include only the final selector packetSelected, to act as an index for that selection sequence in the Selection Sequence Statistics Report Interpretation, which also allows the calculation of the Attained Selection Fraction.

The contiguous Information Elements (dataLinkFrameSection, mplsLabelStackSection, mplsPayloadPacketSection, ipPacketSection, and ipPayloadPacketSection) MAY be encoded with a fixed length field or with a variable sized field. If one of these Information Elements is encoded with a fixed length field whose length is too long for the number of contiguous bytes in the selected packet, padding MUST NOT be used. In this case, the Exporting Process MUST export the information either in a new Template Record with the correct fixed length field, or either in a new Template Record with a variable length field.

Here is an example of a basic Packet Report, with a SelectionSequenceId value of 9 and ipHeaderPacketSection Information Element of 12 bytes, 0x4500 005B A174 0000 FF11 832E, encoded with a fixed length field.

IPFIX Template Record:

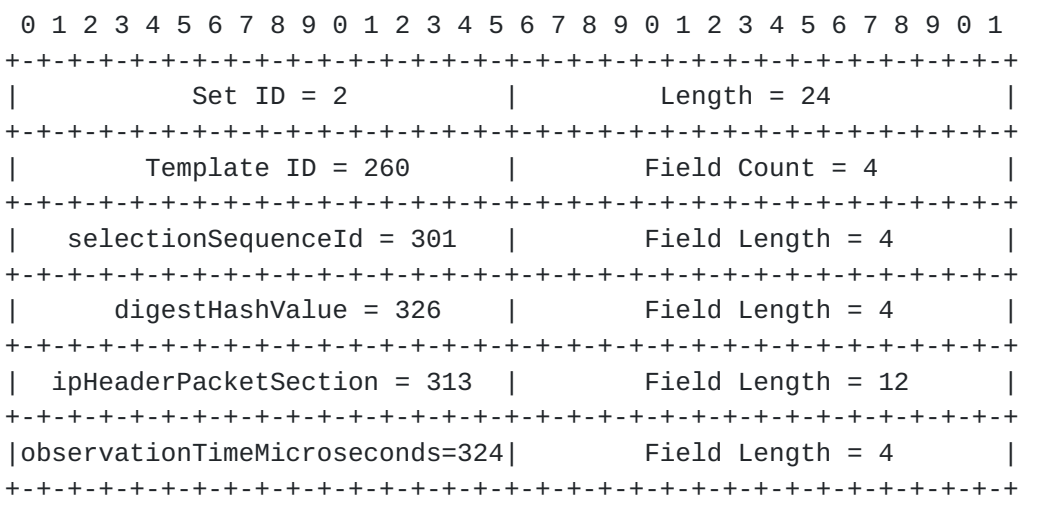

The associated IPFIX Data Record:

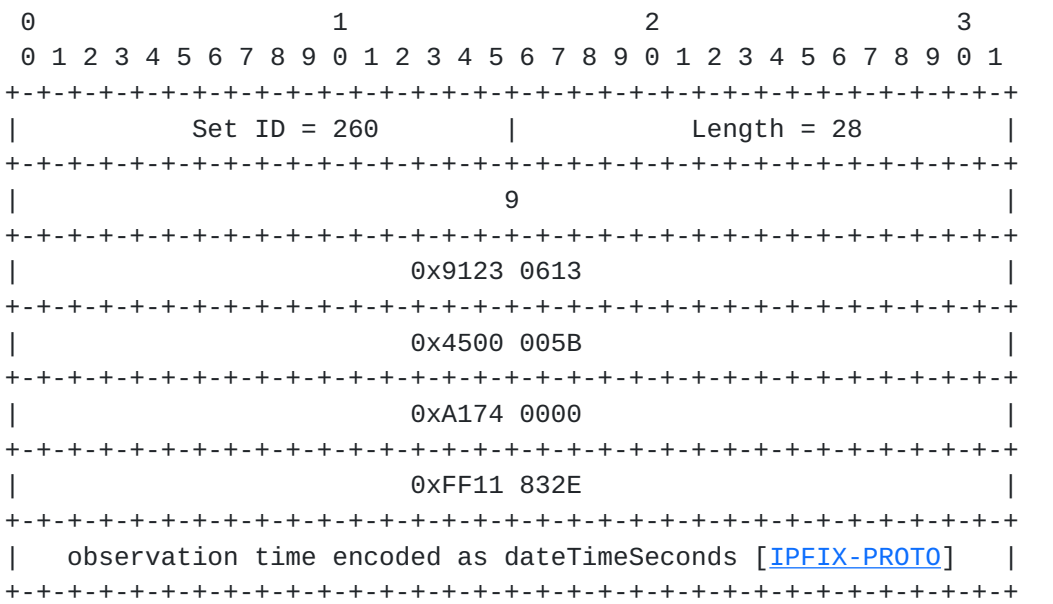

Figure D: Example of a Basic Packet Report

Here is an example of a basic Packet Report, with a SelectionSequenceId value of 9 and ipHeaderPacketSection Information Element of 12 bytes, 0x4500 005B A174 0000 FF11 832E, encoded with a variable sized field.

IPFIX Template Record:

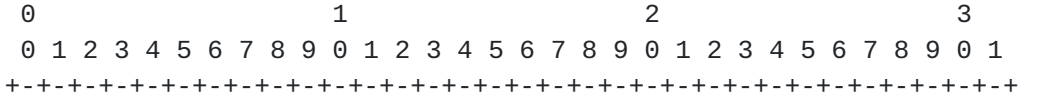

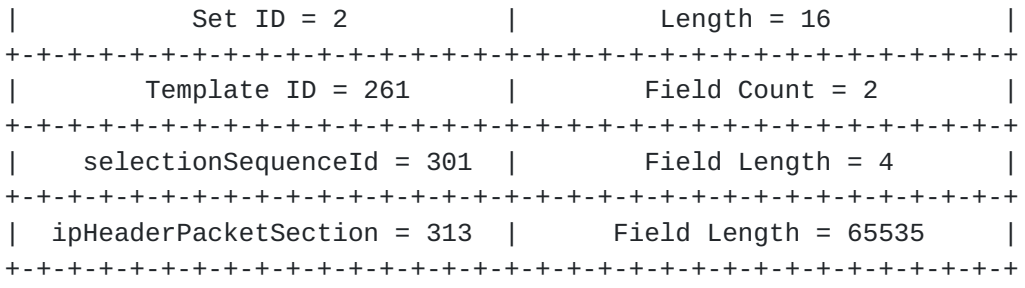

The associated IPFIX Data Record:

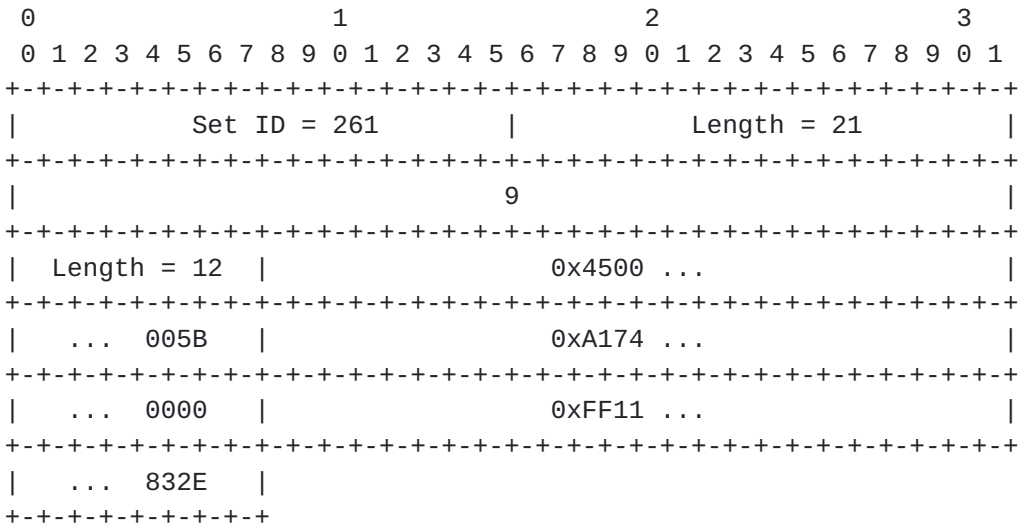

Figure E: Example of a Basic Packet Report,
with a variable sized field

6.4.2 Extended Packet Report

Alternatively to the basic Packet Report, the extended Packet Report MAY contain other Information Elements related to the protocols used in the packet (such as source and destination IP addresses), related to the packet treatment (such as output interface, destination BGP autonomous system [RFC1771]), or related to the Selection State associated with the packet (such as timestamp, hash value).

It is envisaged that selection of fields for extended Packet Reports may be used to reduce reporting bandwidth, in which case the option to report some number of contiguous bytes from the start of the packet, mandatory in the basic Packet Report, may not be exercised. In this case, the Packet Content MAY be omitted. Note this configuration is quite similar to an IPFIX Device for which a Template Record containing information about a single packet is reported.

Example of a detailed Extended Packet Report:

IPFIX Template Record:

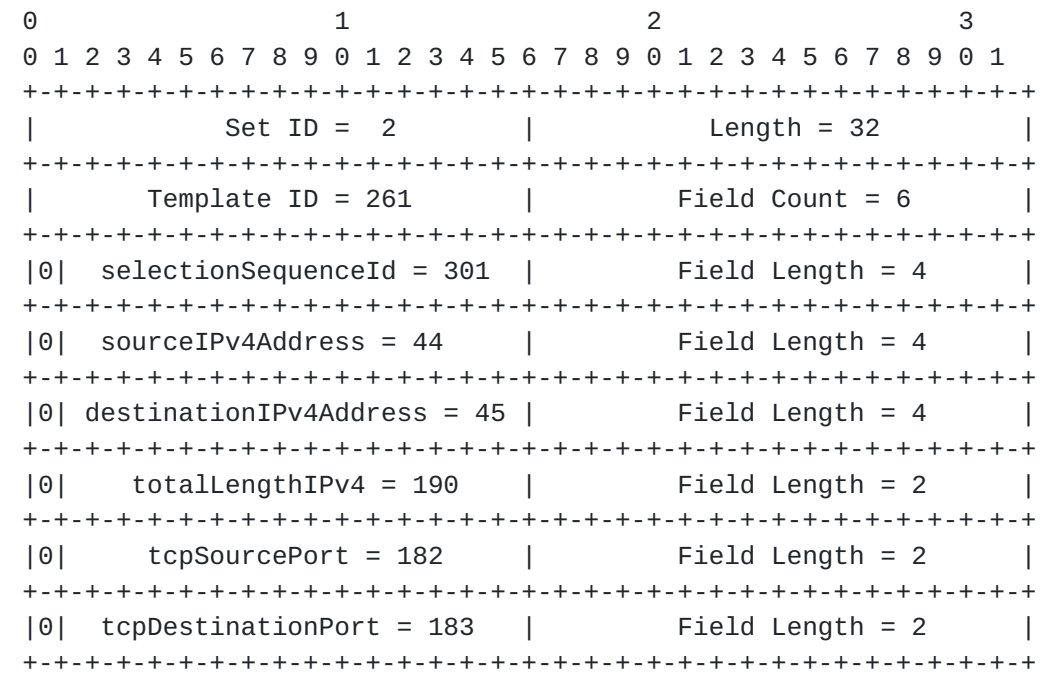

The associated IPFIX Data Record:

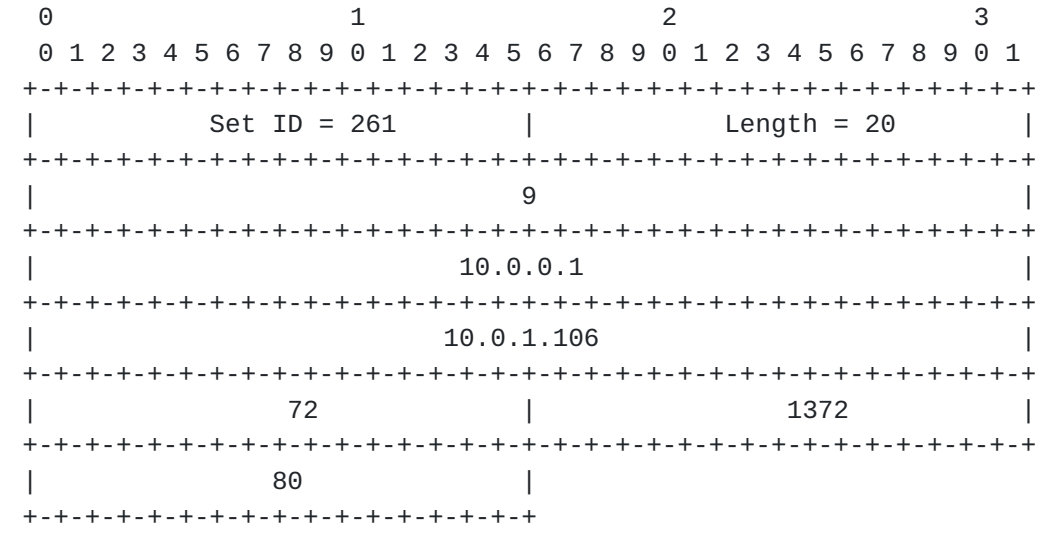

Figure F: Example of an Extended Packet Report

6.5 Report Interpretation

To make full sense of the Packet Reports there are a number of additional pieces of information that must be communicated to the Collector:
- The details about which Selectors and Observation Points are being used within a Selection Sequences MUST be provided using the Selection Sequence Report Interpretation.
- The configuration details of each Selector MUST be provided using the Selector Report Interpretation.
- The Selector ID statistics MUST be provided using the Selection Sequence Statistics Report Interpretation.
- The accuracies of the reported fields MUST be provided using the Accuracy Report Interpretation.

6.5.1 Selection Sequence Report Interpretation

Each Packet Report contains a selectionSequenceId Information Element that identifies the particular combination of Observation Point and Selector(s) used for its selection. For every selectionSequenceId Information Element in use, the PSAMP Device MUST export a Selection Sequence Report Interpretation using an Options Template with the following Information Elements:

```
 Scope:     selectionSequenceId
 Non-Scope: one Information Element representing
                the Observation Point
            selectorId (one or more)
```

An Information Element representing the Observation Point would typically be taken from the ingressInterface, egressInterface, lineCardId, exporterIPv4Address, exporterIPv6Address Information Elements (specified in [IPFIX-INFO]), but not limited to those: any Information Element specified in [IPFIX-INFO] or [PSAMP-INFO] can potentially be used. In case of more complex Observation Points (such as a list of interfaces, a bus, etc..), a new Information Element describing the new type of Observation Point must be specified, along with an option template record describing it in more details (if necessary).

If the packets are selected by a Composite Selector, the Selection Sequence is composed of several Primitive Selectors. In such a case, the Selection Sequence Report Interpretation MUST contain the list of all the Primitive Selector IDs in the Selection Sequence. If multiple Selectors are contained in the Selection Sequence Report Interpretation, the selectorId's MUST be identified in the order they are used.

Example of two Selection Sequences:

```
 Selection Sequence 7 (Filter->Sampling):
   ingressInterface    5
   selectorId          5 (Filter, match IPV4SourceAddress 10.0.0.1)
   selectorId         10 (Sampler, Random 1 out-of ten)

 Selection Sequence 9 (Sampling->Filtering):
   ingressInterface    5
   selectorId         10 (Sampler, Random 1 out-of ten)
   selectorId          5 (Filter, match IPV4SourceAddress 10.0.0.1)
```

IPFIX Options Template Record:

```
 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 = 3           |          Length = 26          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Template ID = 262      |         Field Count = 4       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Scope Field Count =  1    |0|  selectionSequenceId = 301  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       Scope 1 Length = 4      |0|     ingressInterface = 10   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Field Length = 4       |0|      selectorId = 300       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Field Length = 4       |0|      selectorId = 300       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|        Field Length = 4       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The associated IPFIX Data Record:

```
 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 = 262         |           Length = 36         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               7                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               5                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               5                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                              10                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                               9                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

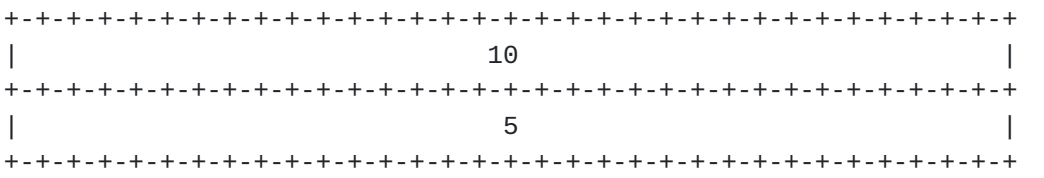

Figure G: Example of a Selection Sequence Report Interpretation

Notes:
* There are two Records here in the same Data Set. Each record defines a different Selection Sequence.
* If, for example, a different Selection Sequence is composed of three Selectors then a different Options Template with three selectorId Information Elements (instead of two) must be used.

6.5.2 Selector Report Interpretation

An IPFIX Data Record, defined by an Option Template Record, MUST be used to send the configuration details of every Selector in use. The Option Template Record MUST contain the selectorId Information Element as the Scope field and the SelectorAlgorithm Information Element followed by some specific configuration parameters:

```
 Scope:     selectorId
 Non-scope: selectorAlgorithm
            algorithm specific Information Elements
```

The algorithm specific Information Elements are specified in the following subsections, depending on the selection method represented by the value of the selectorAlgorithm.

6.5.2.1 Systematic Count-Based Sampling

In systematic count-based Sampling, the start and stop triggers for the Sampling interval are defined in accordance with the spatial packet position (packet count) [PSAMP-TECH].

The REQUIRED algorithm specific Information Elements in the case of systematic count-based Sampling are:

```
   samplingPacketInterval: number of packets selected in a row
   samplingPacketSpace:    number of packets between selections
```

Example of a simple 1 out-of 10 systematic count-based Selector definition, where the samplingPacketInterval is 1 and the samplingPacketSpace is 9.

IPFIX Options Template Record:

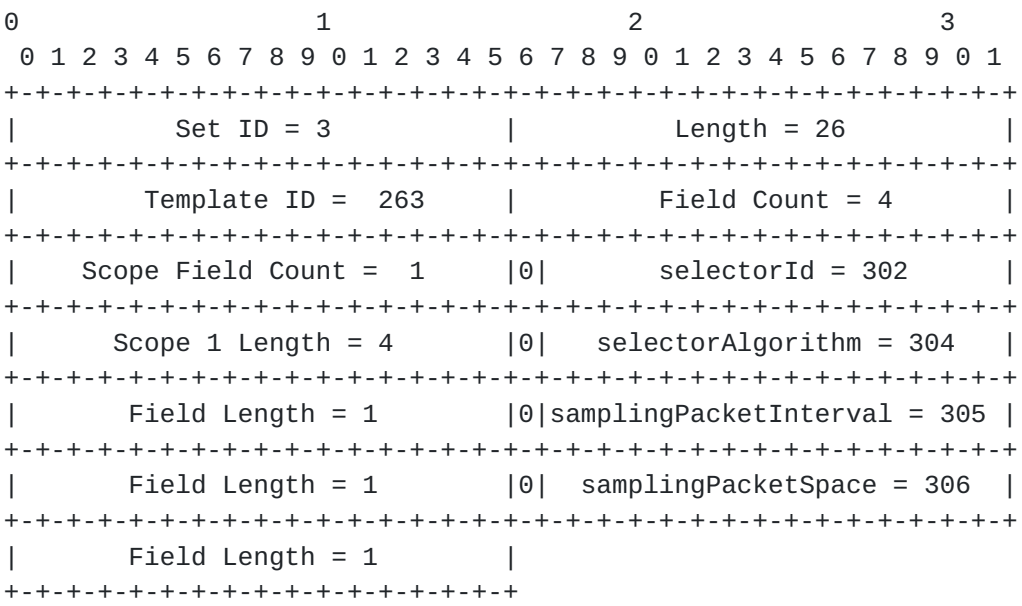

Associated IPFIX Data Record:

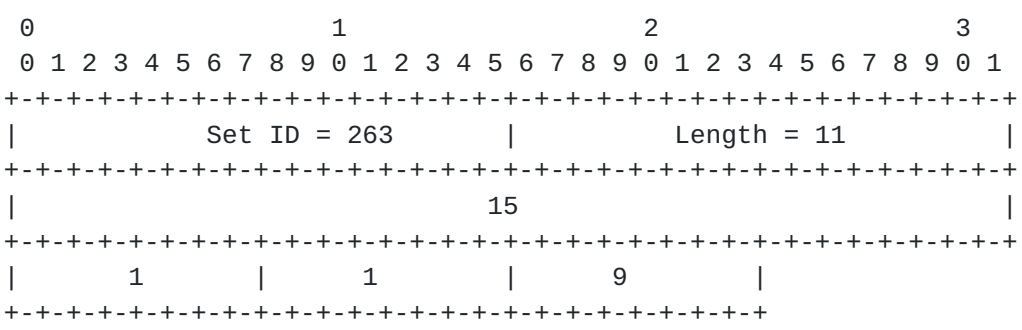

Figure H: Example of the Selector Report Interpretation,
For Systematic Count-Based Sampling

Notes:
* A selectorAlgorithm value of 1 represents systematic count-based Sampling.
* samplingPacketInterval and samplingPacketSpace are of type unsigned32 but are compressed down to one octet here, as allowed by the IPFIX protocol specifications [IPFIX-PROTO].

6.5.2.2 Systematic Time-Based Sampling

In systematic time-based Sampling, the start and stop triggers are used to define the Sampling intervals [PSAMP-TECH]. The REQUIRED

algorithm specific Information Elements in the case of systematic time-based Sampling are:

samplingTimeInterval: time (in us) when packets are selected
samplingTimeSpace: time (in us) between selections

Example of a 100 us out-of 1000 us systematic time-based Selector definition, where the samplingTimeInterval is 100 and the samplingTimeSpace is 900

IPFIX Options Template Record:

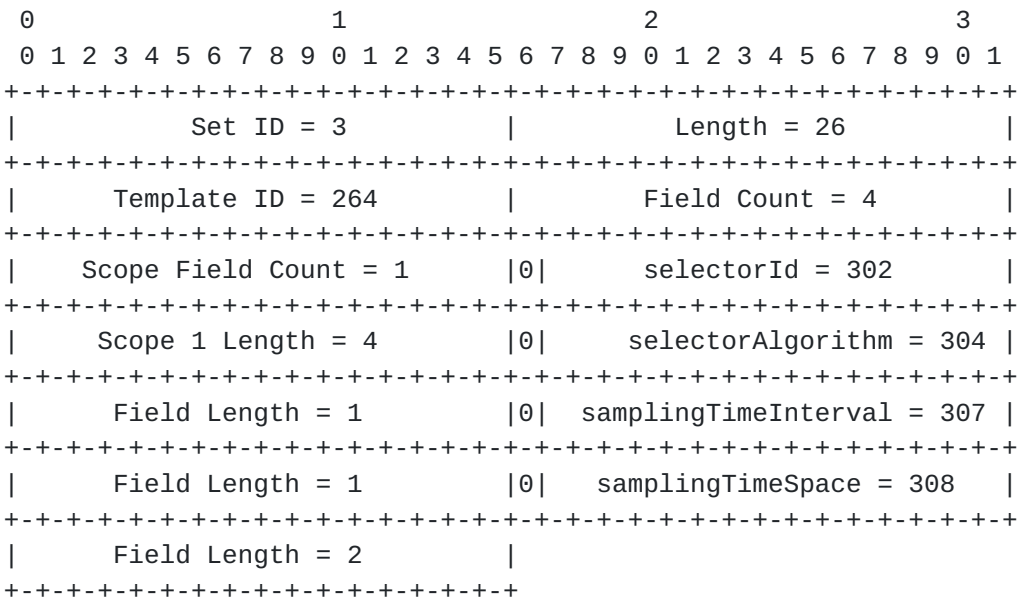

Associated IPFIX Data Record:

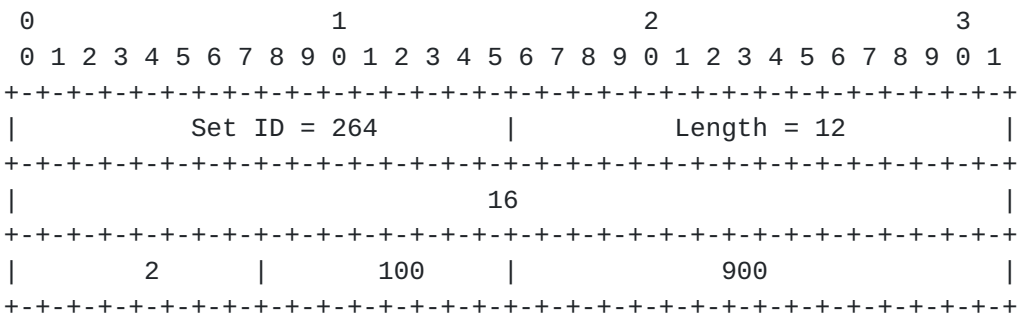

Figure I: Example of the Selector Report Interpretation, For Systematic Time-Based Sampling

Notes:
* A selectorAlgorithm value of 2 represents systematic time-based Sampling.
* samplingTimeInterval and samplingTimeSpace are of type unsigned32

but are compressed down here.

6.5.2.3 Random n-out-of-N Sampling

In random n-out-of-N Sampling, n elements are selected out of the parent population that consists of N elements [PSAMP-TECH]. The REQUIRED algorithm specific Information Elements in case of random n-out-of-N Sampling are:

samplingSize: number of packets selected
samplingPopulation: number of packets in selection population

Example of a 1 out-of 10 random n-out-of-N Sampling Selector:

IPFIX Options Template Record:

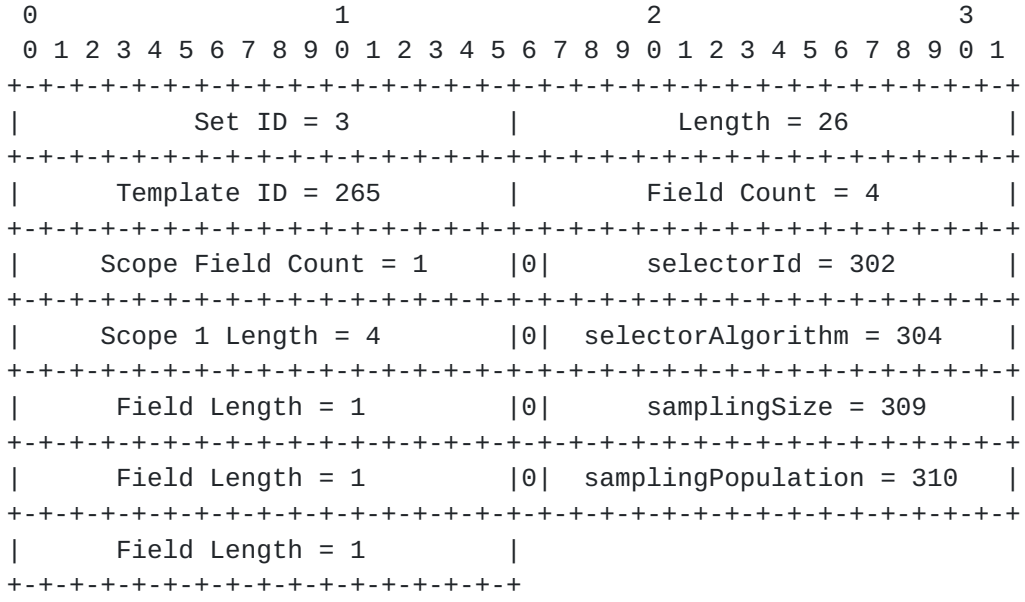

Associated IPFIX Data Record:

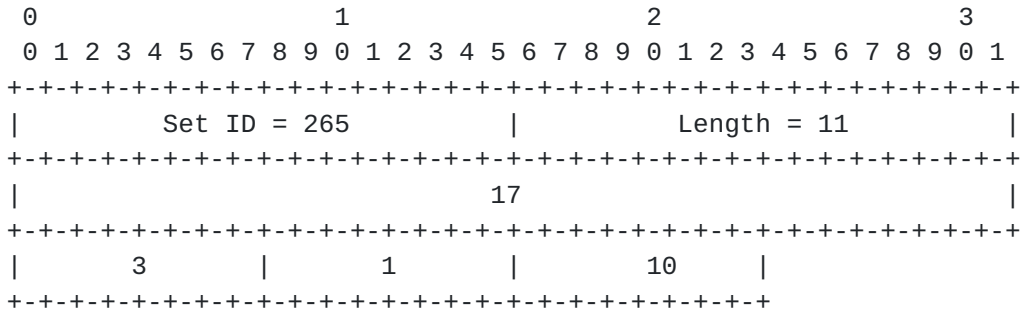

Figure J: Example of the Selector Report Interpretation,

Notes:
* A selectorAlgorithm value of 3 represents Random n-out-of-N Sampling.
* samplingSize and samplingPopulation are of type unsigned32 but are compressed down to one octet here.

6.5.2.4 Uniform Probabilistic Sampling

In uniform probabilistic Sampling, each element has the same probability p of being selected from the parent population [PSAMP-TECH]. The algorithm specific Information Element in case of uniform probabilistic Sampling is:

samplingProbablility: a floating point number for the Sampling probability.

Example of a 15% uniform probability Sampling Selector:

IPFIX Options Template Record:

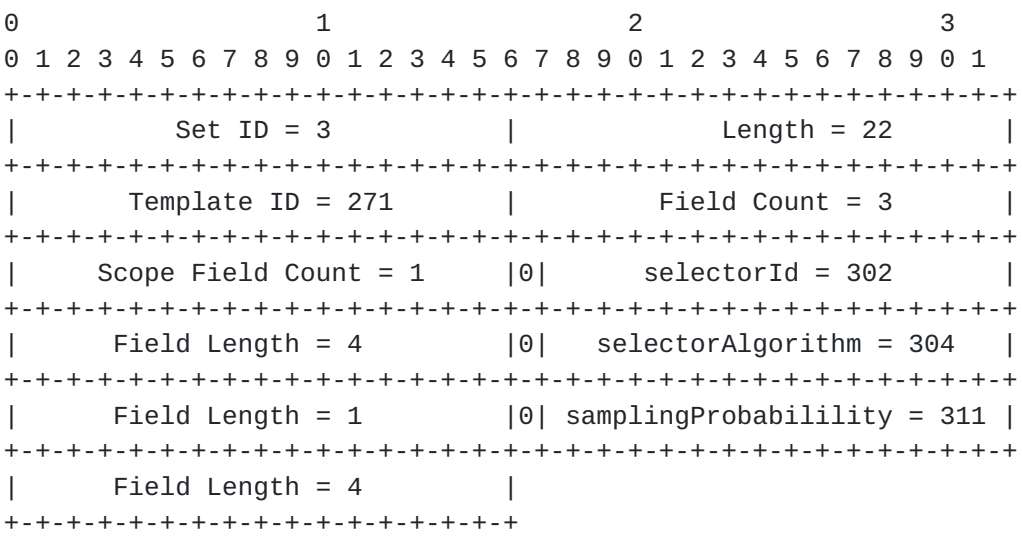

Associated IPFIX Data Record:

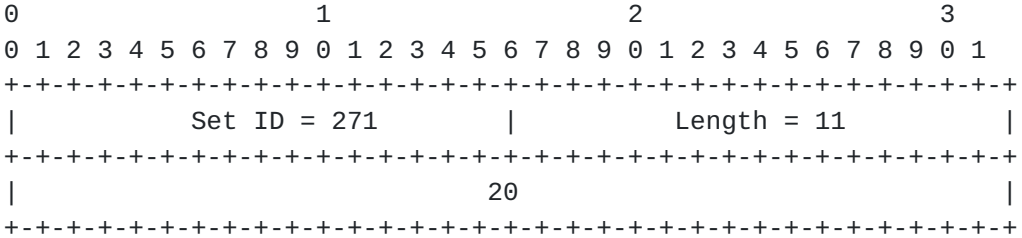

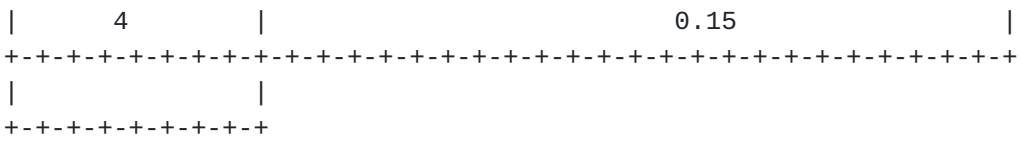

Figure K: Example of the Selector Report Interpretation,
For Uniform Probabilistic Sampling

Notes:
* A selectorAlgorithm value of 4 represents Uniform Probabilistic Sampling.
* samplingProbablility is of type float64 but is compressed down to a float32 here.

6.5.2.5 Property Match Filtering

This classification includes match(es) on field(s) within a packet and/or on properties of the router state. With this method, a packet is selected if a specific field in the packet equals a predefined value.

The algorithm specific Information Elements defining configuration parameters for property match filtering are taken from the full range of available Information Elements.

When multiple different Information Elements are defined, the filter acts as a logical AND. Note that the logical OR is not covered by these PSAMP specifications. The property match Filtering Options Template Record MUST NOT have multiple identical Information Elements. The result of the filter is independent from the order of the Information Elements in the Option Template Record, but the order may be important for implementation purposes, as the first filter will have to work at a higher rate. In any case, an implementation is not constrained to respect the filter ordering as long as the result is the same, and it may even implement the composite Filtering in Filtering in one single step.

Since encryption alters the meaning of encrypted fields, when the Property Match Filtering classification is based on the encrypted field(s) in the packet, it MUST be able to recognize that the field(s) are not available and MUST NOT select those packets unless specifically directed by the Information Element description. Even if they are ignored, the encrypted packets MUST be accounted for in the Selector packetsObserved Information Element [PSAMP-INFO], part of the Selection Sequence Statistics Report Interpretation.

Example of a match based filter Selector, whose rules are:
   IPv4 Source Address = 10.0.0.1
   IPv4 Next-Hop Address = 10.0.1.1

IPFIX Options Template Record:

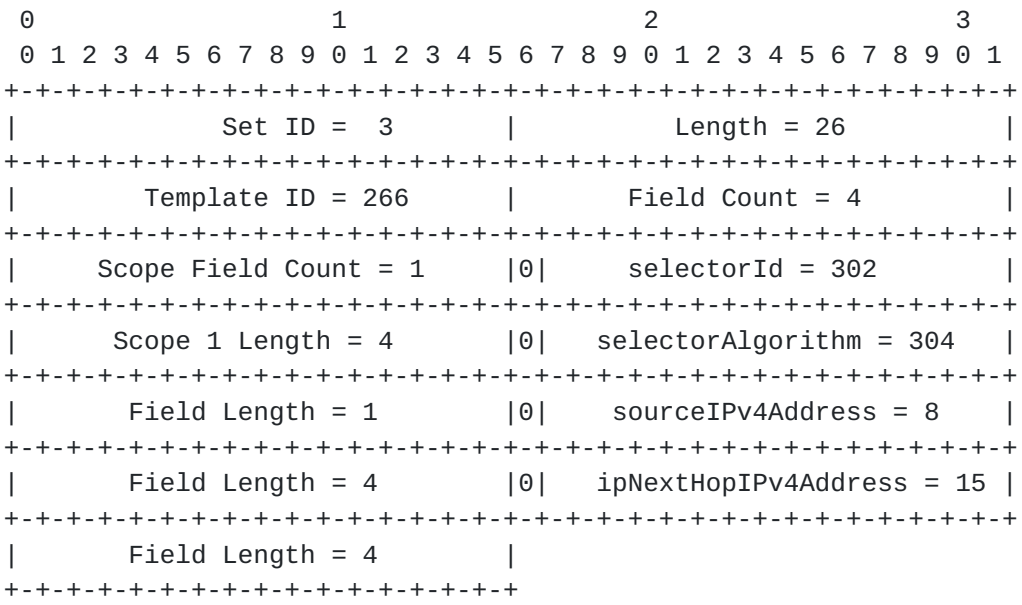

Associated IPFIX Data Record:

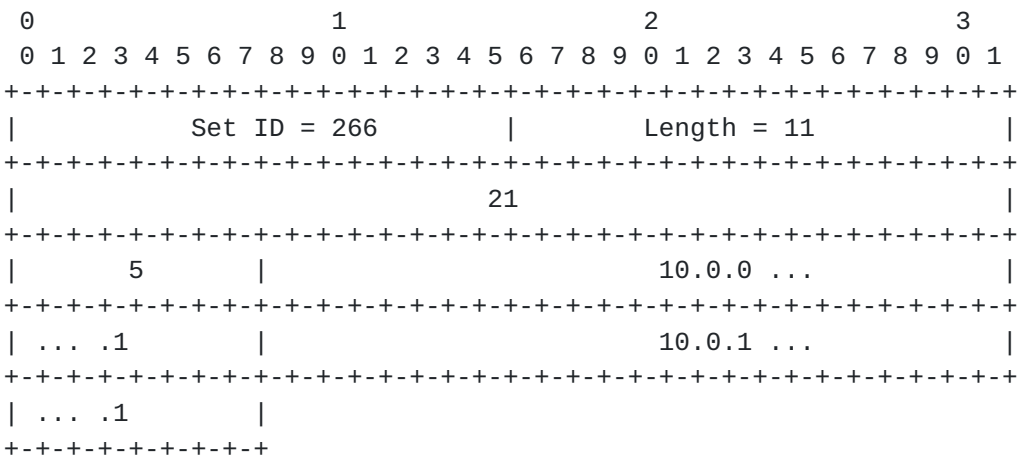

Figure L: Example of the Selector Report Interpretation,
For match based and router state Filtering

Notes:
* A selectorAlgorithm value of 5 represents property match Filtering.
* In this filter there is a mix of information from the packet and information from the router.

6.5.2.6 Hash-Based Filtering

In hash based selection a hash function is run on IPv4 traffic the following fields MUST be used as input to that hash function:

- IP identification field
- Flags field
- Fragment offset
- Source IP address
- Destination IP address
- A number of bytes from the IP payload. The number of bytes and starting offset MUST be configurable if the hash function supports it.

For the bytes taken from the IP payload, IPSX has a fixed offset of 0 bytes and a fixed size of 8 bytes. The number and offset of payload bytes in the BOB function MUST be configurable.

The minimum configuration ranges MUST be as follows:

Number of bytes:	from 8 to 32
Offset:	from 0 to 64

If the selected payload bytes are not available and the hash function can take a variable sized input then the hash function MUST be run with the information which is available and a shorter size. Passing 0 as a substitute for missing payload bytes is only acceptable if the hash function takes a fixed size as is the case with IPSX.

If the hash function can take an initialization value then this value MUST be configurable.

A hash-based selection function MAY be configurable as a digest function. Any selection process which is configured as a digest function MUST have the output value included in the basic packet report for any selected packet.

Each hash function used as a hash-based selector requires its own value for the selectorAlgorithm. Currently we have BOB (6), IPSX (7) and CRC (8) defined and any MAY be used for either Filtering or creating a Packet Digest. Only BOB is recommended though and SHOULD be used.

The REQUIRED algorithm specific Information Elements in case of hash based selection are:

hashIPPayloadOffset	- The payload offset used by a hash based Selector
hashIPPayloadSize	- The payload size used by a hash based Selector

```
hashOutputRangeMin    - One or more values for the beginning of
                        each potential output range.
hashOutputRangeMax    - One or more values for the end of each
                        potential output range.
hashSelectedRangeMin  - One or more values for the beginning of
                        each selected range.
hashSelectedRangeMax  - One or more values for the end of each
                        selected range.
hashDigestOutput      - A boolean value, TRUE if the output from
                        this selector has been configured to be
                        included in the packet report as a packet
                        digest.
```

NOTE: If more than one selection or output range needs to be sent then the minimum and maximum elements may be repeated as needed. These MUST make one or more non-overlapping ranges. The elements SHOULD be sent as pairs of minimum and maximum in ascending order, however if they are sent out of order then there will only be one way to interpret the ranges to produce a non-overlapping range and the Collecting Process MUST be prepared to accept and decode this.

The following algorithm specific Information Element MAY be sent, but is optional for security considerations:

```
hashInitialiserValue  - The initialiser value to the hash function.
```

Since encryption alters the meaning of encrypted fields, when the Hash-Based Filtering classification is based on the encrypted field(s) in the packet, it MUST be able to recognize that the field(s) are not available and MUST NOT select those packets. Even if they are ignored, the encrypted packets MUST be accounted in the Selector packetsObserved Information Element [PSAMP-INFO], part of the Selection Sequence Statistics Report Interpretation.

Example of a hash based filter Selector, whose configuration is:

```
Hash Function           = BOB
Hash IP Payload Offset  = 0
Hash IP Payload Size    = 16
Hash Initialiser Value  = 0x9A3F9A3F
Hash Output Range       = 0 to 0xFFFFFFFF
Hash Selected Range     = 100 to 200 and 400 to 500
```

IPFIX Options Template Record:

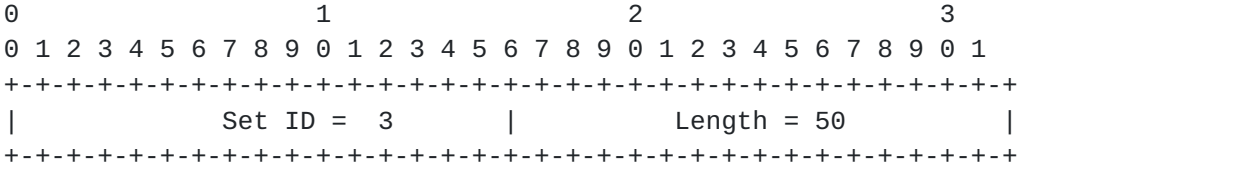

```
|        Template ID = 269      |       Field Count = 8         |
```

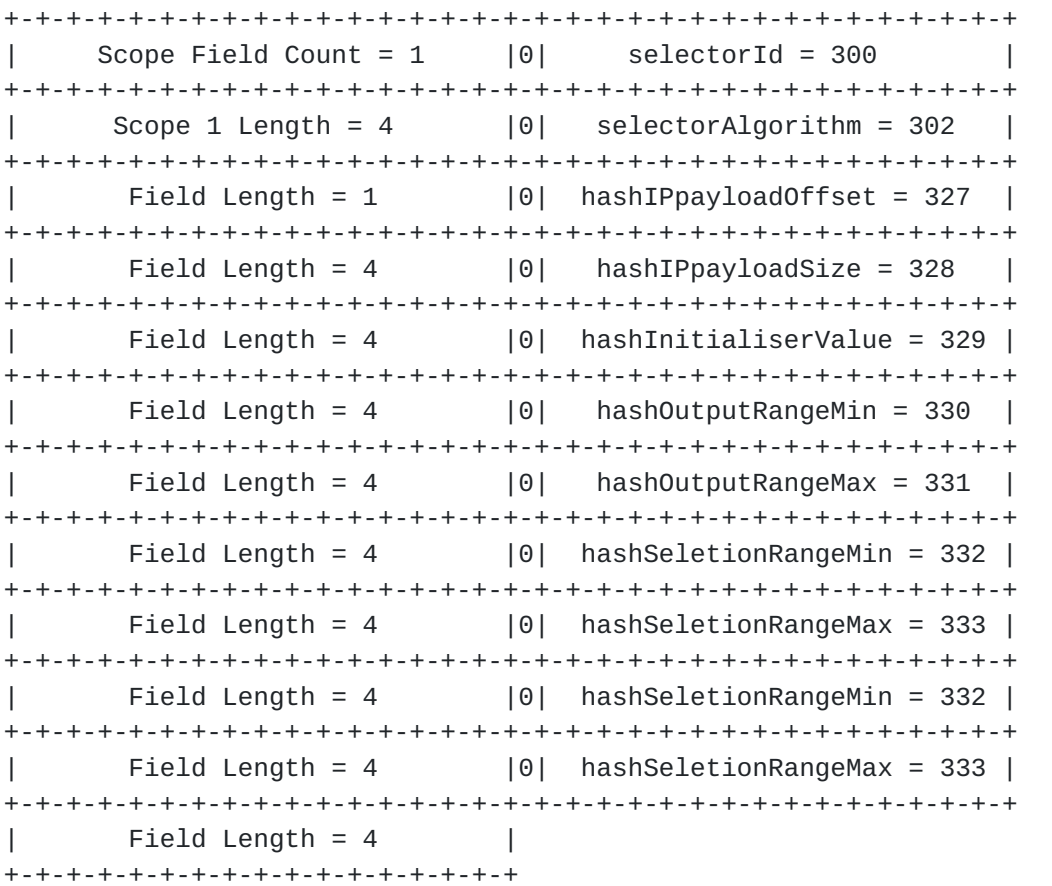

Associated IPFIX Data Record:

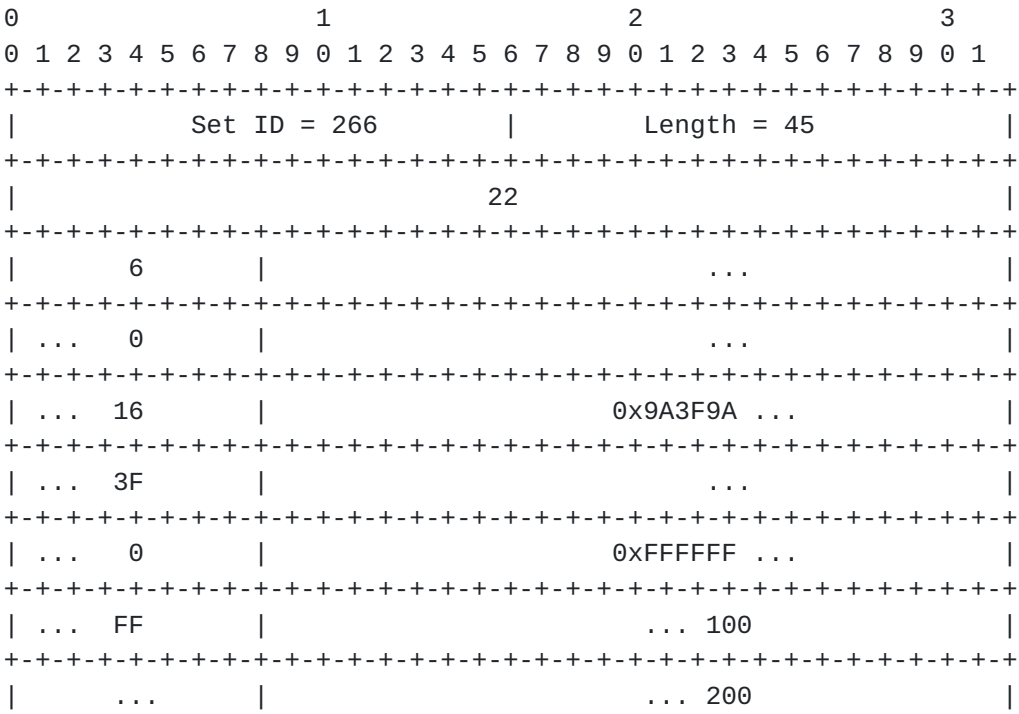

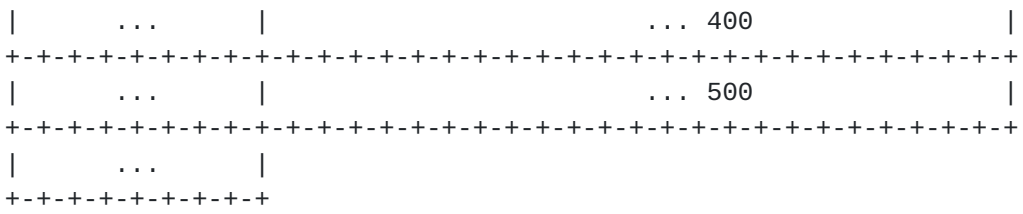

Figure M: Example of the Selector Report Interpretation, for Hash Based Filtering

Notes:
* A selectorAlgorithm value of 6 represents hash-based Filtering using the BOB algorithm.

6.5.2.7 Other Selection Methods

Some potential new selection methods MAY be added. Some of the new selection methods, such as non-uniform probabilistic Sampling and flow state dependent Sampling, are described in [PSAMP-TECH], with further references.

Each new selection method MUST be assigned a unique value for the selectorAlgorithm Information Element. Its configuration parameter(s), along with the way to report it/them with an Options Template, MUST be clearly specified.

6.5.3 Selection Sequence Statistics Report Interpretation

A Selector MAY be used in multiple Selection Sequences. However, each use of a Selector must be independent, so each separate logical instance of a Selector MUST maintain its own individual Selection State and statistics.

The Selection Sequence Statistics Report Interpretation MUST include the number of observed packets (Population Size) and the number of packets selected (Sample Size) by each instance of its Primitive Selectors.

Within a Selection Sequence composed of several Primitive Selectors, the number of packets selected for one Selector is equal to the number of packets seen by the next Selector. The order of the Selectors in the Selection Sequence Statistics Report Interpretation MUST match the order of the Selectors in the Selection Sequence.

If the full set of statistics is not sent as part of the Basic Packet Reports, the PSAMP Device MUST export a Selection Sequence Statistics

Report Interpretation for every Selection Sequence, using an Options Template containing the following Information Elements:

```
 Scope:         selectionSequenceId
 Non-scope:     packetsObserved
                packetsSelected (first)
                ...
                packetsSelected (last)
```

The packetsObserved Information Element [PSAMP-INFO] MUST contain the number of packets seen at the Observation Point, and as a consequence passed to the first Selector in the Selection Sequence. The packetsSelected Information Element [PSAMP-INFO] contains the number of packets selected by a Selector in the Selection Sequence.

The Attained Selection Fraction for the Selection Sequence is calculated by dividing the number of selected packets (packetsSelected Information Element) for the last Selector by the number of observed packets (packetsObserved Information Element). The Attained Selection Fraction can be calculated for each Selector by dividing the number of packets selected for that Selector by the value for the previous Selector.

The statistics for the whole sequence SHOULD be taken at a single logical point in time; the input value for a Selector MUST equal the output value of the previous selector.

The Selection Sequence Statistics Report Interpretation MUST be exported periodically.

Example of Selection Sequence Statistics Report Interpretation:

```
 Selection Sequence 7 (Filter->Sampling):

   Observed   100  (observationPointId  1, Interface 5)
   Selected    50  (selectorId  5, match IPV4SourceAddress 10.0.0.1)
   Selected     6  (selectorId 10, Sampler: Random one out-of ten)

 Selection Sequence 9 (Sampling->Filtering):

   Observed   100  (observationPointId  1, Interface 5)
   Selected    10  (selectorId 10, Sampler: Random one out-of ten)
   Selected     3  (selectorId  5, match IPV4SourceAddress 10.0.0.1)
```

IPFIX Options Template Record:

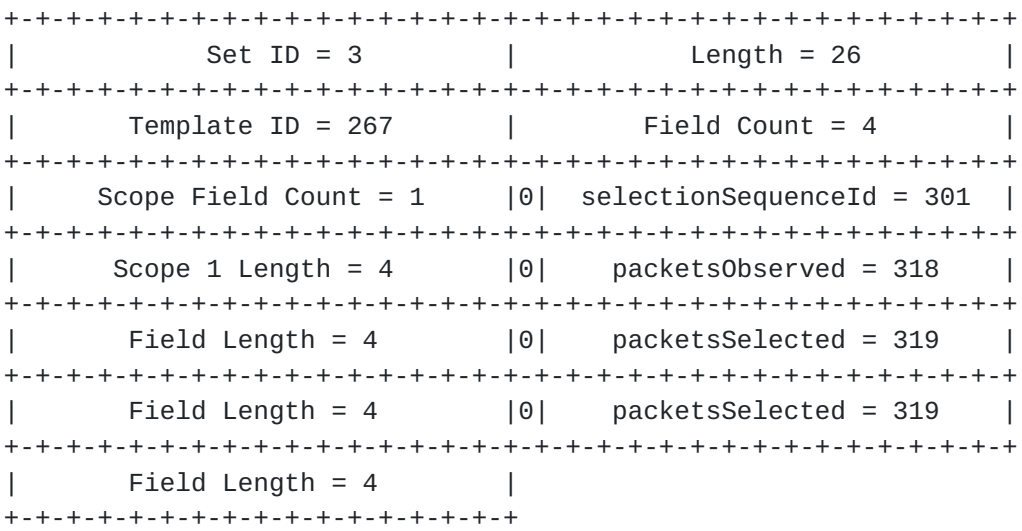

The associated IPFIX Data Record:

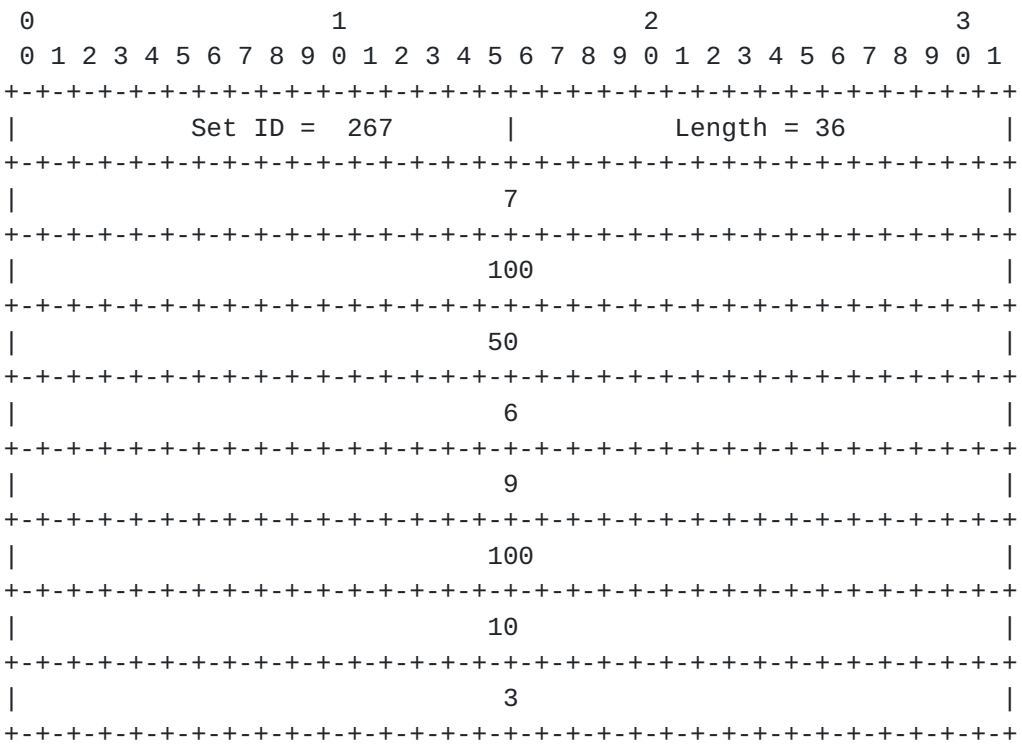

Figure N: Example of the Selection Sequence Statistics Report Interpretation

Notes:
* The Attained Sampling Fractions for Selection Sequence 7 are:
  Filter 10: 50/100
  Sampler 5: 6/50
  Number of samples selected: 6

* The Attained Sampling Fractions for Selection Sequence 9 are:
  Sampler 5: 10/100
  Filter 10: 3/10
  Number of samples selected: 3

6.5.4 Accuracy Report Interpretation

In order for the Collecting Process to determine the inherent accuracy of the reported quantities (for example timestamps), the PSAMP Device SHOULD send an Accuracy Report Interpretation.

The Accuracy Report Interpretation MUST be exported by an Option Template Record with a scope that contains the Information Element for which the accuracy is required. In case the accuracy is specific to a template, a second scope containing the templateId value MUST be added to the Option Template Record. The accuracy SHOULD be reported either with the fixedError Information Element [PSAMP-INFO], or with the relativeError Information Element [PSAMP-INFO].

Accuracy Report Interpretation using the fixedError Information Element:
 Scope: informationElementId
 Non-scope: fixedError

Accuracy Report Interpretation using the fixedError Information Element and a double scope:
 Scope: templateId
 informationElementId
 Non-scope: fixedError

Accuracy Report Interpretation using the relativeError Information Element:
 Scope: informationElementId
 Non-scope: relativeError

Accuracy Report Interpretation using the relativeError Information Element and a double scope:
 Scope: templateId
 informationElementId
 Non-scope: relativeError

For example, the accuracy of an Information Element whose Abstract Data Type is dateTimeMilliseconds [IPFIX-INFO], for which the unit is specified as milliseconds, can be specified with the fixedError Information Element with the milliseconds units. In this case, the error interval is the Information Element value +/- the value reported in the fixedError.

For example, the accuracy of an Information Element to estimate the accuracy of a sampled flow, for which the unit would be specified in octets, can be specified with the relativeError Information Element with the octet units. In this case, the error interval is the Information Element value +/- the value reported in the relativeError times the reported Information Element value.

Alternatively to reporting either the fixedError Information Element or the relativeError Information Element in the Accuracy Report Interpretation, both Information Elements MAY be present. This scenario could help in more complex situations where the system clock drifts, on the top of having its own accuracy, during the duration of a measurement.

If the accuracy of a reported quantity changes on the Metering Process, a new Accuracy Report Interpretation MUST be generated. The Collecting Process MUST keep the accuracy of the latest Accuracy Report Interpretation.

Example of an Accuracy Report Interpretation using the fixedError Information Element and a double scope: the timeMicroseconds contained in the Template 5 has an accuracy of +/- 2 ms, represented by the fixedError Information Element.

```
 Scope:     templateId = 6
            informationElementId = timeMicroseconds
 Non-scope: fixedError = 2 ms
```

IPFIX Options Template Record:

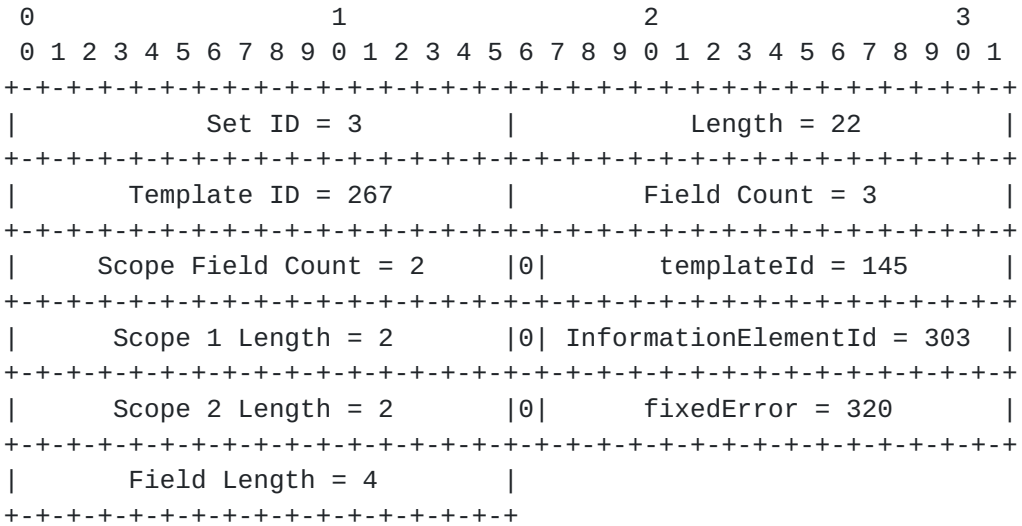

The associated IPFIX Data Record:

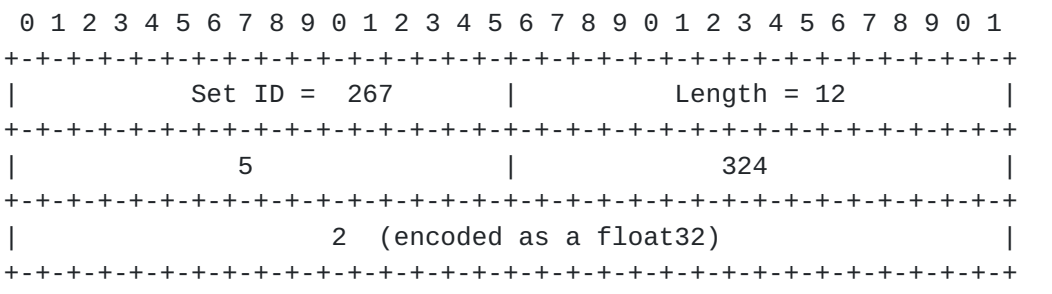

Figure O: Example of the Selection Sequence Statistics Report Interpretation

Notes:
* fixedError is of type float64 but is compressed down to a float32 here.

The second example displays an Accuracy Report Interpretation using the relativeError Information Element and a single scope: the timeMicroseconds has an error of 5 percents, represented by the proportionalAccuracy Information Element.
Scope: informationElementId = timeMicroseconds
Non-scope: relativeError = 0.05

IPFIX Options Template Record:

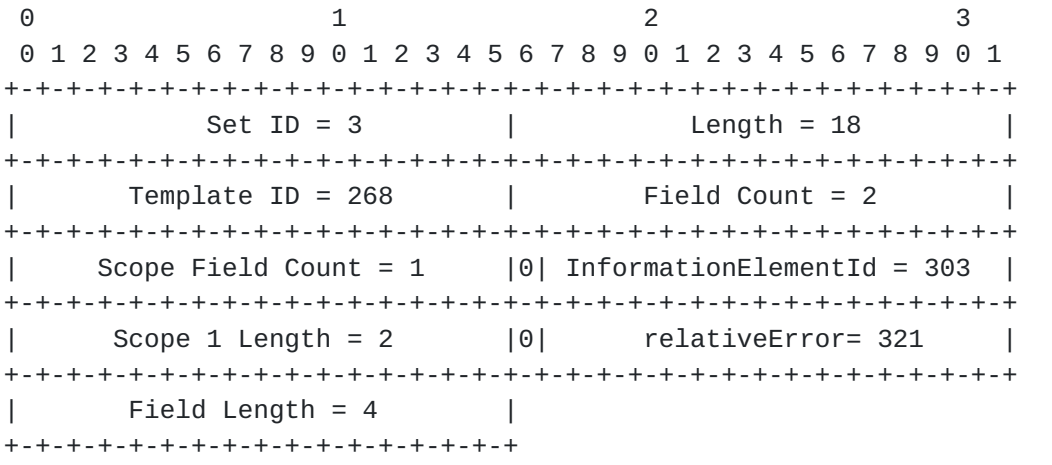

The associated IPFIX Data Record:

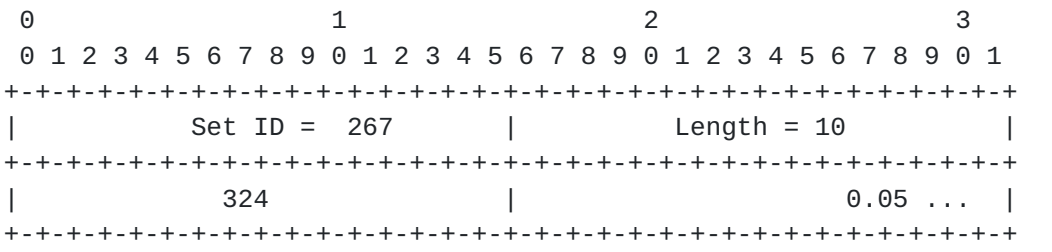

```
  | ...(encoded as a float32)     |
```

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure P: Example of the Selection Sequence Statistics Report Interpretation

Notes:
* relativeError is of type float64 but is compressed down to a float32 here.

7. Security Considerations

As IPFIX has been selected as the PSAMP export protocol and as the PSAMP security requirements are not stricter than the IPFIX security requirements, refer to the IPFIX export protocol [IPFIX-PROTO] for the security considerations.

In the basic Packet Report, a PSAMP Device exports some number of contiguous bytes from the start of the packet, including the packet header (which includes link layer, network layer and other encapsulation headers) and some subsequent bytes of the packet payload. The PSAMP Device SHOULD NOT export the full payload of conversations, as this would mean wiretapping [RFC2804]. The PSAMP Device MUST respect local privacy laws.

8. IANA Considerations

The PSAMP Protocol, as set out in this document, has two sets of assigned numbers. Considerations for assigning them are discussed in this section, using the example policies as set out in the "Guidelines for IANA Considerations" document IANA-RFC [RFC2434].

8.1 IPFIX Related Considerations

As the PSAMP protocol uses the IPFIX protocol, refer to the IANA considerations section in [IPFIX-PROTO] for the assignments of numbers used in the protocol and for the numbers used in the information model.

8.2 PSAMP Related Considerations

Each new selection method MUST be assigned a unique value for the selectorAlgorithm Information Element. Its configuration parameter(s), along with the way to report it/them with an Options Template, MUST be clearly specified.

New assignments for the PSAMP selection method will be administered by IANA, on a First Come First Served basis [RFC2434], subject to

Expert Review [RFC2434], i.e. review by one of a group of experts designated by an IETF Operations and Management Area Director. The group of experts must double check the Information Elements definitions with already defined Information Elements for completeness, accuracy and redundancy. Those experts will initially be drawn from the Working Group Chairs and document editors of the IPFIX and PSAMP Working Groups.

9. References

9.1 Normative References

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

[RFC2434] H. Alvestrand, T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2434, October 1998

[PSAMP-TECH] T. Zseby, M. Molina, N. Duffield, S. Niccolini, F. Raspall, "Sampling and Filtering Techniques for IP Packet Selection" draft-ietf-psamp-sample-tech-07.txt

[PSAMP-INFO] T. Dietz, F. Dressler, G. Carle, B. Claise, "Information Model for Packet Sampling Exports", draft-ietf-psamp-info-03.txt

[IPFIX-ARCH] G. Sadasivan, N. Brownlee, B. Claise, J. Quittek, "Architecture Model for IP Flow Information Export" draft-ietf-ipfix-arch-12.txt"

[IPFIX-INFO] J. Quittek, S. Bryant, B. Claise, J. Meyer, "Information Model for IP Flow Information Export" draft-ietf-ipfix-info-13.txt

[IPFIX-PROTO] B. Claise (Editor) "Specification of the IPFIX Protocol for the Exchange of IP Traffic Flow Information", draft-ietf-ipfix-protocol-23.txt

9.2 Informative References

[PSAMP-FMWK] D. Chiou, B. Claise, N. Duffield, A. Greenberg, M. Grossglauser, P. Marimuthu, J. Rexford, G. Sadasivan, "A Framework for Passive Packet Measurement" draft-ietf-psamp-framework-10.txt

[RFC1771] Y. Rekhter, T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995

[RFC1889] Schulzrinne, H., Casner, S., Frederick, R., Jacobson, V., "RTP: A Transport Protocol for Real-Time Applications", RFC 1889, January 1996

[RFC2804] IAB, IESG, "IETF Policy on Wiretapping", RFC 2804, May 2000

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

10. Acknowledgments

The authors would like to thank the PSAMP group, especially Paul Aitken for fruitful discussions and for proofreading the document several times.

Authors' Addresses

Benoit Claise
Cisco Systems
De Kleetlaan 6a b1
1831 Diegem
Belgium
Phone: +32 2 704 5622
E-mail: bclaise@cisco.com

Juergen Quittek
NEC Europe Ltd.
Network Laboratories
Kurfuersten-Anlage 36
69115 Heidelberg
Germany
Phone: +49 6221 90511-15
Email: quittek@ccrle.nec.de

Andrew Johnson
Cisco Systems
96 Commercial Quay
Edinburgh EH6 6LX
Scotland
Phone: +44 131 561 3641
Email: andrjohn@cisco.com

Intellectual Property Statement

on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society