Network Working Group
Internet-Draft
Expires: April 26, 2007

T. Dietz
NEC Europe Ltd.
F. Dressler
University of Erlangen-Nuremberg
G. Carle
University of Tuebingen
B. Claise
P. Aitken
Cisco Systems

October 23, 2006

Information Model for Packet Sampling Exports <draft-ietf-psamp-info-05.txt>

Status of this Memo

Copyright Notice

Abstract

This memo defines an information model for the Packet Sampling (PSAMP) protocol. It is used by the PSAMP protocol for encoding

sampled packet data and information related to the sampling process. As the PSAMP protocol is based on the IPFIX protocol, this information model is an extension to the IPFIX information model.

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

Packet sampling techniques are required for various measurement scenarios. The packet sampling (PSAMP) protocol provides mechanisms for packet selection using different filtering and sampling techniques. A standard way for the export and storage of such sampled packet data is required. The definition of the PSAMP information and data model is based on the IP Flow Information eXport (IPFIX) protocol [I-D.ietf-ipfix-protocol]. The PSAMP protocol document [I-D.ietf-psamp-protocol] describes how to use the IPFIX protocol in the PSAMP context.

This document examines the IPFIX information model [I-D.ietf-ipfix-info] and extends it to meet the PSAMP requirements. Therefore, the structure of this document is strongly based on the IPFIX document. It complements the PSAMP protocol specification by providing an appropriate PSAMP information model. The main part of this document, section 8, defines the list of Information Elements to be transmitted by the PSAMP protocol. Sections 6 and 5 describe the data types and Information Element properties used within this document and their relationship to the IPFIX information model.

The main body of section 8 was generated from a XML document. The XML-based specification of the PSAMP Information Elements can be used for automatically checking syntactical correctness of the specification. Furthermore it can be used - in combination with the IPFIX information model - for automated code generation. The resulting code can be used in PSAMP protocol implementations to deal with processing PSAMP information elements.

For that reason, the XML document that served as source for section 8 is attached to this document in Appendix A.

Note that although partially generated from the attached XML documents, the main body of this document is normative while the appendices are informational.

2. PSAMP Documents Overview

[I-D.ietf-psamp-framework]: "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.

[I-D.ietf-psamp-sample-tech]: "Sampling and Filtering Techniques for IP Packet Selection", describes the set of packet selection techniques supported by PSAMP.

[I-D.ietf-psamp-protocol]: "Packet Sampling (PSAMP) Protocol Specifications" specifies the export of packet information from a PSAMP Exporting Process to a PSAMP Collecting Process.

[I-D.ietf-psamp-info]: "Information Model for Packet Sampling Exports" (this document), defines an information and data model for PSAMP.

3. Relationship between PSAMP and IPFIX

As described in the PSAMP protocol draft [I-D.ietf-psamp-protocol] a PSAMP data record can be seen as a very special IPFIX Data Record. It represents an IPFIX flow containing only a single packet. Therefore, the IPFIX information model can be used as a basis for PSAMP reports.

Nevertheless, there are properties required in PSAMP reports which cannot be modelled using the current IPFIX information model. This document describes extensions to the IPFIX model which allow the modelling of information and data required by PSAMP.

Some of these extensions allow the export of what may be considered sensitive information. Refer to the Security Considerations section for a fuller discussion.

Note that the export of sampled data may not need all the information elements defined by the IPFIX information model [I-D.ietf-ipfix-info], as discussed in sections 6.2 and 6.3 of the PSAMP Framework [I-D.ietf-psamp-framework].

4. Terminology

4.1 General Terminology

- IETF: The Internet Engineering Task Force

  http://www.ietf.org

- IPFIX: The IETF IP Flow Information eXport working group

  http://www.ietf.org/html.charters/ipfix-charter.html

- PSAMP: The IETF Packet SAMPling working group

  http://www.ietf.org/html.charters/psamp-charter.html

- IANA: Internet Assigned Numbers Authority

  http://www.iana.org

- RFC: Requests for Comments

  http://www.rfc-editor.org

- ISO: International Organisation for Stadardisation

  http://www.iso.org

- IEC: International Electrotechnical Commission

  http://www.iec.ch

4.2 PSAMP Terminology

The relevant PSAMP terminology has been copied from [I-D.ietf-psamp-sample-tech] into this document.

- Observed Packet Stream

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

- Selection Process

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

- Packet Stream

  A packet stream denotes a subset of the Observed Packet Stream that flows past some specified point within the Selection Process.

- 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.

- Selector

  A Selector defines the action of a Selection Process on a single packet of its input.

- 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.

4.3 IPFIX Terminology

As the IPFIX export protocol is used to export the PSAMP information, the relevant IPFIX terminology from [I-D.ietf-ipfix-protocol] is copied over in this document.

- Observation Point

  An Observation Point is a location in the network where packets can be observed.

- Observation Domain

  An Observation Domain is the largest set of Observation Points for which Flow information can be aggregated by a Metering Process.

- 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.

- Flow Record

  A Flow Record contains information about a specific Flow that was observed at an Observation Point.

- Metering Process

  The Metering Process generates Flow Records.

- Exporting Process

  The Exporting Process sends Flow Records to one or more Collecting Processes. The Flow Records are generated by one or more Metering Processes.

- IPFIX Device

  An IPFIX Device hosts at least one Exporting Process. It may host further Exporting processes and arbitrary numbers of Observation Points and Metering Process.

- Collecting Process

  A Collecting Process receives Flow Records from one or more Exporting Processes.

- 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.

- Information Element

  An Information Element is a protocol and encoding independent description of an attribute which may appear in an IPFIX Record.

5. Properties of a PSAMP Information Element

The PSAMP Information Elements are in accordance with the definitions of IPFIX. Therefore we do not repeat the properties in this draft. Refer to sections 2.1 through 2.3 of the IPFIX Information Model [I-D.ietf-ipfix-info]. Nevertheless, we strongly recommend to define the optional "units" property for every information element (if applicable).

6. Type Space

The PSAMP Information Elements MUST be constructed from the basic data types described in section 3 of the IPFIX Information Model [I-D.ietf-ipfix-info]. To avoid duplicated work and to keep consistency between IPFIX and PSAMP, the data types are not repeated in this document.

7. Overloading Information Elements

Information Elements won't be overloaded with multiple meanings or re-used for multiple purposes. Different Information Elements will be allocated for each requirement.

In particular, special information will be encoded in new Information Elements as necessary, and not be encoded in the selection method.

Although the prescence of certain other Information Elements allows the selection method to be inferred, a separate Information Element is provided for the selectorAlgorithm, e.g. for including in scope info and depicting the contents of composites.

8. The PSAMP Information Elements

This section describes the Information Elements used by the PSAMP exporting functions.

Each Information Element defined in section 8.2 below is allocated a unique identifier in accordance with section 4 of the IPFIX information model [I-D.ietf-ipfix-info]. The assignments are controlled by IANA as an extension of the IPFIX Information Model.

The Information Elements described by the IPFIX information model [I-D.ietf-ipfix-info] are used by the PSAMP export functions where applicable. To avoid inconsistencies between the IPFIX and the PSAMP information and data models, only those Information Elements that are not already described by the IPFIX information model are defined here.

8.1 PSAMP Usage of IPFIX Attributes

Some Information Elements defined by the IPFIX information model are not needed by the PSAMP protocol.

This section lists additional Information Elements that are needed in the PSAMP context and introduces their usage.

List of additional PSAMP Information Elements:

- 300 - observationPointId
- 301 - selectionSequenceId
- 302 - selectorId
- 303 - informationElementId
- 304 - selectorAlgorithm
- 305 - samplingPacketInterval
- 306 - samplingPacketSpace
- 307 - samplingTimeInterval
- 308 - samplingTimeSpace
- 309 - samplingSize
- 310 - samplingPopulation
- 311 - samplingProbability
- 312 - dataLinkFrameSize
- 313 - ipHeaderPacketSection
- 314 - ipPayloadPacketSection
- 315 - dataLinkFrameSection
- 316 - mplsLabelStackSection
- 317 - mplsPayloadPacketSection
- 318 - packetsObserved
- 319 - packetsSelected
- 320 - fixedError
- 321 - relativeError
- 322 - observationTimeSeconds
- 323 - observationTimeMilliseconds
- 324 - observationTimeMicroseconds
- 325 - observationTimeNanoseconds
- 326 - digestHashValue
- 327 - hashIPPayloadOffset
- 328 - hashIPPayloadSize
- 329 - hashOutputRangeMin
- 330 - hashOutputRangeMax
- 331 - hashSelectedRangeMin
- 332 - hashSelectedRangeMax
- 333 - hashDigestOutput
- 334 - hashInitialiserValue

8.2 Additional PSAMP Information Elements

8.2.1 observationPointId

Description:
ID of the Observation Point. Unique in the Observation Domain.
Abstract Data Type: unsigned64
ElementId: 300

Status: current

8.2.2 selectionSequenceId

Description:
From all the packets observed at an Observation Point, a subset of packets is selected by a sequence of one or more Selectors. The selectionSequenceId is a unique value per Observation Domain, describing the Observation Point and the sequence of Selectors through which the packets are selected.
Abstract Data Type: unsigned64
ElementId: 301
Status: current

8.2.3 selectorId

Description:
The Selector ID is the unique ID identifying a Primitive Selector. Each Primitive Selector must have a unique ID in the Observation Domain.
Abstract Data Type: unsigned16
Data Type Semantics: identifier
ElementId: 302
Status: current

8.2.4 informationElementId

Description:
Contains the ID of another Information Element.
Abstract Data Type: unsigned16
Data Type Semantics: identifier
ElementId: 303
Status: current

8.2.5 selectorAlgorithm

Description:
Specifies the Selector algorithm (e.g., filter, sampler, hash) that was used on a packet.

The following Selector algorithms are currently defined:

* 1 Systematic count-based sampling
* 2 Systematic time-based sampling
* 3 Random n-out-of-N sampling
* 4 Uniform probabilistic sampling

* 5 Property match filtering
* 6 Hash based filtering using BOB
* 7 Hash based filtering using IPSX
* 8 Hash based filtering using CRC

The parameters for most of these algorithms are defined in this information model. Some parameters for these algorithms are not covered by this information model since they very much depend on the underlying hardware.

In future, this list will be maintained by IANA. IANA can update this information element as long as there's a new RFC specifying the algorithm and any new Information Elements which are required.

Abstract Data Type: unsigned16
Data Type Semantics: identifier
ElementId: 304
Status: current

8.2.6 samplingPacketInterval

Description:
Number of packets that are consecutively sampled. For example a value of 100 means that 100 contiguous packets are sampled.

This information element is used to describe the configuration of a systematic count-based sampling Selector.

Abstract Data Type: unsigned32
ElementId: 305
Status: current
Units: packets

8.2.7 samplingPacketSpace

Description:
The number of packets between two "samplingPacketInterval"s. A value of 100 means that the next interval starts after 100 packets (which are not sampled) when the current "samplingPacketInterval" is over.

This information element is used to describe the configuration of a systematic count-based sampling Selector.

Abstract Data Type: unsigned32
ElementId: 306
Status: current

Units: packets

8.2.8 samplingTimeInterval

Description:
Time interval in microseconds in which all arriving packets are sampled.

This information element is used to describe the configuration of a systematic time-based sampling Selector.
Abstract Data Type: dateTimeMicroseconds
ElementId: 307
Status: current
Units: microseconds

8.2.9 samplingTimeSpace

Description:
The time interval in microseconds between two "samplingTimeInterval"s. A value of 100 means that the next interval starts after 100 microseconds (in which no packets are sampled) when the current "samplingTimeInterval" is over.

This information element is used to describe the configuration of a systematic time-based sampling Selector.
Abstract Data Type: dateTimeMicroseconds
ElementId: 308
Status: current
Units: microseconds

8.2.10 samplingSize

Description:
The number of elements taken from the parent Population for random sampling algorithms.

This information element is used to describe the configuration of a random n-out-of-N sampling Selector.
Abstract Data Type: unsigned32
ElementId: 309
Status: current
Units: packets

8.2.11 samplingPopulation

Description:
The number of elements in the parent Population for random sampling algorithms.

This information element is used to describe the configuration of a random n-out-of-N sampling Selector.
Abstract Data Type: unsigned32
ElementId: 310
Status: current
Units: packets

8.2.12 samplingProbability

Description:
The probability that a packet is sampled, expressed as a value between 0 and 1. The probability is equal for every packet. A value of 0 means no packet was sampled since the probability is 0.

This information element is used to describe the configuration of a uniform probabilistic sampling Selector.
Abstract Data Type: float64
ElementId: 311
Status: current

8.2.13 dataLinkFrameSize

Description:
The size of the data link frame.

The data link layer is defined in [ISO/IEC.7498-1:1994].
Abstract Data Type: unsigned32
ElementId: 312
Status: current

8.2.14 ipHeaderPacketSection

Description:
This information element carries a series of octets from the start of the IP header of a sampled packet.

With sufficient length, this element also reports octets from the IP payload, subject to [RFC2804] and the PSAMP WG charter. See the Security Considerations section.

The size of the exported section may be constrained due to limitations in the IPFIX protocol.

If insufficient octets are available for the length specified in the Template, the Information Element MUST NOT be padded.

Abstract Data Type: variable length octetArray
ElementId: 313
Status: current

8.2.15 ipPayloadPacketSection

Description:
This information element carries a series of octets from the start of the IP payload of a sampled packet.

The IPv4 payload is that part of the packet which follows the IPv4 header and any options, which [RFC0791] refers to as "data" or "data octets". e.g., see the examples in [RFC0791] APPENDIX A.

The IPv6 payload is the rest of the packet following the 40 octet IPv6 header. Note that any extension headers present are considered part of the payload. See [RFC2460] for the IPv6 specification.

The size of the exported section may be constrained due to limitations in the IPFIX protocol.

If insufficient octets are available for the length specified in the Template, the Information Element MUST NOT be padded.

Abstract Data Type: variable length octetArray
ElementId: 314
Status: current

8.2.16 dataLinkFrameSection

Description:
This information element carries the first n octets from the data link frame of a sampled packet.

The data link layer is defined in [ISO/IEC.7498-1:1994].

The size of the exported section may be constrained due to limitations in the IPFIX protocol.

If insufficient octets are available for the length specified in the Template, the Information Element MUST NOT be padded.

Abstract Data Type: variable length octetArray
ElementId: 315
Status: current

8.2.17 mplsLabelStackSection

Description:
This information element carries the first n octets from the MPLS label stack of a sampled packet.

With sufficient length, this element also reports octets from the MPLS payload, subject to [RFC2804] and the PSAMP WG charter. See the Security Considerations section.

See [RFC3031] for the specification of MPLS packets.
See [RFC3032] for the specification of the MPLS label stack.

The size of the exported section may be constrained due to limitations in the IPFIX protocol.

If insufficient octets are available for the length specified in the Template, the Information Element MUST NOT be padded.

Abstract Data Type: variable length octetArray
ElementId: 316
Status: current

8.2.18 mplsPayloadPacketSection

Description:
This information element carries the first n octets from the MPLS payload of a sampled packet, being data that follows immediately after the MPLS label stack.

See [RFC3031] for the specification of MPLS packets.
See [RFC3032] for the specification of the MPLS label stack.

The size of the exported section may be constrained due to limitations in the IPFIX protocol.

If insufficient octets are available for the length specified in the Template, the Information Element MUST NOT be padded.

Abstract Data Type: variable length octetArray

ElementId: 317
Status: current

8.2.19 packetsObserved

Description:
Number of packets observed by a Selector.
Abstract Data Type: unsigned64
ElementId: 318
Status: current
Units: packets

8.2.20 packetsSelected

Description:
Number of packets selected by a Selector.
Abstract Data Type: unsigned64
ElementId: 319
Status: current
Units: packets

8.2.21 fixedError

Description:
Specifies the maximum possible positive or negative error interval
of the reported value for a given Information Element.
Abstract Data Type: float64
ElementId: 320
Status: current
Units: The units of the Information Element for which the error is
specified.

8.2.22 relativeError

Description:
Specifies the maximum possible positive or negative error ratio
for a given Information Element.
Abstract Data Type: float64
ElementId: 321
Status: current

8.2.23 observationTimeSeconds

Description:

The absolute time of an observation.
Abstract Data Type: dateTimeSeconds
ElementId: 322
Status: current
Units: seconds

8.2.24 observationTimeMilliseconds

Description:
The absolute time of an observation.
Abstract Data Type: dateTimeMilliseconds
ElementId: 323
Status: current
Units: milliseconds

8.2.25 observationTimeMicroseconds

Description:
The absolute time of an observation.
Abstract Data Type: dateTimeMicroseconds
ElementId: 324
Status: current
Units: microseconds

8.2.26 observationTimeNanoseconds

Description:
The absolute time of an observation.
Abstract Data Type: dateTimeNanoseconds
ElementId: 325
Status: current
Units: nanoseconds

8.2.27 digestHashValue

Description:
The value from the digest hash function.
Abstract Data Type: unsigned64
ElementId: 326
Status: current

8.2.28 hashIPPayloadOffset

Description:

The IP payload offset used by a hash based Selector.
Abstract Data Type: unsigned64
ElementId: 327
Status: current

8.2.29 hashIPPayloadSize

Description:
The IP payload size used by a hash based Selector.
Abstract Data Type: unsigned64
ElementId: 328
Status: current

8.2.30 hashOutputRangeMin

Description:
A value for the beginning of a hash function's potential output range.
Abstract Data Type: unsigned64
ElementId: 329
Status: current

8.2.31 hashOutputRangeMax

Description:
A value for the end of a hash function's potential output range.
Abstract Data Type: unsigned64
ElementId: 330
Status: current

8.2.32 hashSelectedRangeMin

Description:
A value for the beginning of a hash function's selected range.
Abstract Data Type: unsigned64
ElementId: 331
Status: current

8.2.33 hashSelectedRangeMax

Description:
A value for the end of a hash function's selected range.
Abstract Data Type: unsigned64
ElementId: 332

Status: current

8.2.34 hashDigestOutput

Description:
A boolean value, TRUE if the output from this hash Selector has been configured to be included in the packet report as a packet digest, else FALSE.
Abstract Data Type: boolean
ElementId: 333
Status: current

8.2.35 hashInitialiserValue

Description:
The initialiser value to the hash function.
Abstract Data Type: unsigned64
ElementId: 334
Status: current

9. Security Considerations

The PSAMP information model itself does not directly introduce security issues. Rather it defines a set of attributes which may for privacy or business issues be considered sensitive information.

Specifically, the Information Elements pertaining to packet sections MUST target no more than the packet header, some subsequent bytes of the packet, and encapsulating headers if present. Full packet capture of arbitrary packet streams is explicitly out of scope, per [RFC2804] and the PSAMP WG charter.

The underlying protocol used to exchange the information described here must therefore apply appropriate procedures to guarantee the integrity and confidentiality of the exported information. Such protocols are defined in separate documents, specifically the IPFIX protocol document [I-D.ietf-ipfix-protocol].

10. IANA Considerations

This document defines an initial set of PSAMP Information Elements as specified in [I-D.ietf-psamp-sample-tech], as an extension to the IPFIX Information Elements [I-D.ietf-ipfix-info]. New assignments for PSAMP Information Elements will be administered according to rules explained in the "IANA Consideration" section of the IPFIX Information Model document [I-D.ietf-ipfix-info].

Note that the PSAMP Information Element IDs were initially started at

the value 300, in order to leave a gap for any ongoing IPFIX work requiring new Information Elements. It is expected that this gap in the Information Element numbering will be filled in by IANA with new IPFIX Information Elements.

Appendix B defines an XML schema which may be used to create consistent machine readable extensions to the IPFIX information model. This schema introduces a new namespace, which will be assigned by IANA according to [RFC3688].

In future the selectorAlgorithm registry will be maintained by IANA. IANA can update this information element as long as there's a new RFC specifying the algorithm and any new Information Elements which are required.

11. References

11.1 Normative References

[I-D.ietf-psamp-sample-tech]
Zseby, T., Molina, M., Duffield, N., Niccolini, S., and F. Raspall, "Sampling and Filtering Techniques for IP Packet Selection", draft-ietf-psamp-sample-tech-07 (work in progress), July 2005.

[I-D.ietf-psamp-protocol]
Claise, B., Quittek, J., and A. Johnson, "Packet Sampling (PSAMP) Protocol Specifications", draft-ietf-psamp-protocol-03 (work in progress), December 2005.

[I-D.ietf-ipfix-info]
Quittek, J., Bryant, S., Claise, B., and J. Meyer, "Information Model for IP Flow Information Export", draft-ietf-ipfix-info-11 (work in progress), September 2005.

[I-D.ietf-ipfix-protocol]
Claise, B., "Specification of the IPFIX Protocol for the Exchange of IP Traffic Flow Information", draft-ietf-ipfix-protocol-23 (work in progress), October 2006.

[ISO/IEC.7498-1:1994]
International Organization for Standardization, "Information technology -- Open Systems Interconnection -- Basic Reference Model: The Basic Mode", ISO Standard 7498-1:1994, June 1996.

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

11.2 Informative References

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

[I-D.ietf-ipfix-architecture] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", draft-ietf-ipfix-architecture-09 (work in progress), August 2005.

[I-D.ietf-psamp-framework] Duffield, N., Chiou, D., Claise, B., Greenberg, A., Grossglauser, M., Marimuthu, P., Rexford, J., and G. Sadasivan, "A Framework for Packet Selection and Reporting", draft-ietf-psamp-framework-10 (work in progress), January 2005.

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

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3032] Rosen, E., Tappan, D., Rekhter, Y., Fedorkow, G., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.

[RFC3444] Pras, A. and J. Schoenwaelder, "On the Difference between Information Models and Data Models", RFC 3444, January 2003.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

[RFC3470] Hollenbeck, S., Rose, M., and L. Masinter, "Guidelines for the Use of Extensible Markup Language (XML) within IETF Protocols", BCP 70, RFC 3470, January 2003.

[RFC3688] Mealling, M., "The IETF XML Registry", BCP 81, RFC 3688, January 2004.

Authors' Addresses

Thomas Dietz
NEC Europe Ltd.
Network Laboratories
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Phone: +49 6221 90511-28
Email: dietz@netlab.nec.de
URI: http://www.netlab.nec.de/

Falko Dressler
University of Erlangen-Nuremberg
Dept. of Computer Sciences
Martensstr. 3
Erlangen 91058
Germany

Phone: +49 9131 85-27914
Email: dressler@informatik.uni-erlangen.de
URI: http://www7.informatik.uni-erlangen.de/~dressler

Georg Carle
University of Tuebingen
Wilhelm-Schickard-Institute for Computer Science
Auf der Morgenstelle 10C
Tuebingen 71076
Germany

Phone: +49 7071 29-70505
Email: carle@informatik.uni-tuebingen.de
URI: http://net.informatik.uni-tuebingen.de/~carle/

Benoit Claise
Cisco Systems
De Kleetlaan 6a b1
Degem 1813
Belgium

Phone: +32 2 704 5622
Email: bclaise@cisco.com

Paul Aitken
Cisco Systems
96 Commercial Quay
Edinburgh EH6 6LX
Scotland

Phone: +44 131 561 3616
Email: paitken@cisco.com
URI: http://www.cisco.com/

Appendix A. Formal Specification of PSAMP Information Elements

This appendix contains a formal description of the PSAMP information model XML document. Note that this appendix is of informational nature, while the text in section Section 8 generated from this appendix is normative.

Using a formal and machine readable syntax for the information model enables the creation of PSAMP aware tools which can automatically adapt to extensions to the information model, by simply reading updated information model specifications.

The wide availability of XML aware tools and libraries for client devices is a primary consideration for this choice. In particular libraries for parsing XML documents are readily available. Also mechanisms such as the Extensible Stylesheet Language (XSL) allow for transforming a source XML document into other documents. This draft was authored in XML and transformed according to [RFC2629].

It should be noted that the use of XML in exporters, collectors or other tools is not mandatory for the deployment of PSAMP. In particular, exporting processes do not produce or consume XML as part of their operation. It is expected that PSAMP collectors MAY take advantage of the machine readability of the information model vs. hardcoding their behavior or inventing proprietary means for accommodating extensions.

Using XML-based specifications does not currently address possible

IANA implications associated with XML Namespace URIs. The use of Namespaces as an extension mechanism implies that an IANA registered Namespace URI should be available and that directory names below this base URI be assigned for relevant IETF specifications. The authors are not aware of this mechanism today.

```
<?xml version="1.0" encoding="UTF-8"?>
<fieldDefinitions><!-- xmlns="http://www.ietf.org/ipfix"-->
  <field name="observationPointId" dataType="unsigned64"
         fieldId="300" status="current" group="common">
    <description>
      <paragraph>
        ID of the Observation Point.
        Unique in the Observation Domain.
      </paragraph>
    </description>
  </field>

  <field name="selectionSequenceId" dataType="unsigned64"
         fieldId="301" status="current" group="common">
    <description>
      <paragraph>
        From all the packets observed at an Observation Point, a
        subset of packets is selected by a sequence of one or more
        Selectors.  The selectionSequenceId is a unique value per
        Observation Domain, describing the Observation Point and the
        sequence of Selectors through which the packets are selected.
      </paragraph>
    </description>
  </field>

  <field name="selectorId" dataType="unsigned16"
         dataTypeSemantics="identifier"
         fieldId="302" status="current" group="common">
    <description>
      <paragraph>
          The Selector ID is the unique ID identifying a Primitive
          Selector. Each Primitive Selector must have a unique ID
          in the Observation Domain.
      </paragraph>
    </description>
  </field>

  <field name="informationElementId" dataType="unsigned16"
         dataTypeSemantics="identifier"
         fieldId="303" status="current" group="common">
```

```
      <description>
        <paragraph>
          Contains the ID of another Information Element.
        </paragraph>
      </description>
    </field>

    <field name="selectorAlgorithm" dataType="unsigned16"
           dataTypeSemantics="identifier"
           fieldId="304" status="current" group="common">
      <description>
        <paragraph>
          Specifies the Selector algorithm (e.g., filter, sampler,
          hash) that was used on a packet.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The following Selector algorithms are currently defined:
        </paragraph>

        <vspace blankLines="1" />

        <itemlist>
          <item>1 Systematic count-based sampling</item>
          <item>2 Systematic time-based sampling</item>
          <item>3 Random n-out-of-N sampling</item>
          <item>4 Uniform probabilistic sampling</item>
          <item>5 Property match filtering</item>
          <item>6 Hash based filtering using BOB</item>
          <item>7 Hash based filtering using IPSX</item>
          <item>8 Hash based filtering using CRC</item>
        </itemlist>

        <vspace blankLines="1" />

        <paragraph>
          The parameters for most of these algorithms
          are defined in this information model. Some parameters for
          these algorithms are not covered by this information model
          since they very much depend on the underlying hardware.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          In future, this list will be maintained by IANA.
```

```
          IANA can update this information element as long as
          there's a new RFC specifying the algorithm and
          any new Information Elements which are required.
        </paragraph>
      </description>
    </field>

    <field name="samplingPacketInterval" dataType="unsigned32"
           fieldId="305" status="current" group="common">
      <description>
        <paragraph>
          Number of packets that are consecutively sampled.
          For example a value of 100 means that 100 contiguous
          packets are sampled.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          This information element is used to describe the
          configuration of a systematic count-based sampling Selector.
        </paragraph>
      </description>
      <units>packets</units>
    </field>

    <field name="samplingPacketSpace" dataType="unsigned32"
           fieldId="306" status="current" group="common">
      <description>
        <paragraph>
          The number of packets between two
          "samplingPacketInterval"s. A value of 100 means that the
          next interval starts after 100 packets (which are not
          sampled) when the current "samplingPacketInterval" is over.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          This information element is used to describe the
          configuration of a systematic count-based sampling Selector.
        </paragraph>
      </description>
      <units>packets</units>
    </field>

    <field name="samplingTimeInterval" dataType="dateTimeMicroseconds"
           fieldId="307" status="current" group="common">
```

```
      <description>
        <paragraph>
          Time interval in microseconds in which all arriving
          packets are sampled.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          This information element is used to describe the
          configuration of a systematic time-based sampling Selector.
        </paragraph>
      </description>
      <units>microseconds</units>
    </field>

    <field name="samplingTimeSpace" dataType="dateTimeMicroseconds"
           fieldId="308" status="current" group="common">
      <description>
        <paragraph>
          The time interval in microseconds between two
          "samplingTimeInterval"s. A value of 100 means that the
          next interval starts after 100 microseconds (in which no
          packets are sampled) when the current "samplingTimeInterval"
          is over.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          This information element is used to describe the
          configuration of a systematic time-based sampling Selector.
        </paragraph>
      </description>
      <units>microseconds</units>
    </field>

    <field name="samplingSize" dataType="unsigned32"
           fieldId="309" status="current" group="common">
      <description>
        <paragraph>
          The number of elements taken from the parent
          Population for random sampling algorithms.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
```

```
          This information element is used to describe the
          configuration of a random n-out-of-N sampling Selector.
        </paragraph>
      </description>
      <units>packets</units>
    </field>

    <field name="samplingPopulation" dataType="unsigned32"
           fieldId="310" status="current" group="common">
      <description>
        <paragraph>
          The number of elements in the parent Population
          for random sampling algorithms.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          This information element is used to describe the
          configuration of a random n-out-of-N sampling Selector.
        </paragraph>
      </description>
      <units>packets</units>
    </field>

    <field name="samplingProbability" dataType="float64"
           fieldId="311" status="current" group="common">
      <description>
        <paragraph>
          The probability that a packet is sampled,
          expressed as a value between 0 and 1.
          The probability is equal for every packet.
          A value of 0 means no packet was sampled
          since the probability is 0.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          This information element is used to describe the
          configuration of a uniform probabilistic sampling Selector.
        </paragraph>
      </description>
    </field>

    <field name="dataLinkFrameSize"
           dataType="unsigned32"
           fieldId="312" status="current" group="common">
```

```
      <description>
        <paragraph>
          The size of the data link frame.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The data link layer is defined in
          <xref target="ISO/IEC.7498-1:1994"/>.
        </paragraph>
      </description>
    </field>

    <field name="ipHeaderPacketSection"
           dataType="variable length octetArray"
           fieldId="313" status="current" group="common">
      <description>
        <paragraph>
          This information element carries a series of octets
          from the start of the IP header of a sampled packet.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          With sufficient length, this element also reports
          octets from the IP payload, subject to [RFC2804] and
          the PSAMP WG charter. See the Security Considerations
          section.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The size of the exported section may be constrained
          due to limitations in the IPFIX protocol.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          If insufficient octets are available for the length specified
          in the Template, the Information Element MUST NOT be padded.
        </paragraph>
      </description>
    </field>
```

```
    <field name="ipPayloadPacketSection"
           dataType="variable length octetArray"
           fieldId="314" status="current" group="common">
      <description>
        <paragraph>
          This information element carries a series of octets
          from the start of the IP payload of a sampled packet.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The IPv4 payload is that part of the packet which follows the
          IPv4 header and any options, which <xref target="RFC0791"/>
          refers to as "data" or "data octets".
          e.g., see the examples in <xref target="RFC0791"/>
          APPENDIX A.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The IPv6 payload is the rest of the packet following the
          40 octet IPv6 header. Note that any extension headers present
          are considered part of the payload.
          See <xref target="RFC2460"/> for the IPv6 specification.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The size of the exported section may be constrained
          due to limitations in the IPFIX protocol.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          If insufficient octets are available for the length specified
          in the Template, the Information Element MUST NOT be padded.
        </paragraph>
      </description>
    </field>

    <field name="dataLinkFrameSection"
           dataType="variable length octetArray"
           fieldId="315" status="current" group="common">
      <description>
```

```
        <paragraph>
          This information element carries the first n octets
          from the data link frame of a sampled packet.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The data link layer is defined in
          <xref target="ISO/IEC.7498-1:1994"/>.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The size of the exported section may be constrained
          due to limitations in the IPFIX protocol.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          If insufficient octets are available for the length specified
          in the Template, the Information Element MUST NOT be padded.
        </paragraph>
      </description>
    </field>

    <field name="mplsLabelStackSection"
           dataType="variable length octetArray"
           fieldId="316" status="current" group="common">
      <description>
        <paragraph>
          This information element carries the first n octets
          from the MPLS label stack of a sampled packet.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          With sufficient length, this element also reports
          octets from the MPLS payload, subject to [RFC2804] and
          the PSAMP WG charter. See the Security Considerations
          section.
        </paragraph>

        <vspace blankLines="1" />
```

```
        <paragraph>
          See <xref target="RFC3031"/>
          for the specification of MPLS packets.
          <vspace blankLines="0" />
          See <xref target="RFC3032"/>
          for the specification of the MPLS label stack.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The size of the exported section may be constrained
          due to limitations in the IPFIX protocol.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          If insufficient octets are available for the length specified
          in the Template, the Information Element MUST NOT be padded.
        </paragraph>
      </description>
    </field>

    <field name="mplsPayloadPacketSection"
           dataType="variable length octetArray"
           fieldId="317" status="current" group="common">
      <description>
        <paragraph>
          This information element carries the first n octets
          from the MPLS payload of a sampled packet, being data
          that follows immediately after the MPLS label stack.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          See <xref target="RFC3031"/>
          for the specification of MPLS packets.
          <vspace blankLines="0" />
          See <xref target="RFC3032"/>
          for the specification of the MPLS label stack.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          The size of the exported section may be constrained
```

```
          due to limitations in the IPFIX protocol.
        </paragraph>

        <vspace blankLines="1" />

        <paragraph>
          If insufficient octets are available for the length specified
          in the Template, the Information Element MUST NOT be padded.
        </paragraph>
      </description>
    </field>

    <field name="packetsObserved" dataType="unsigned64"
           fieldId="318" status="current" group="common">
      <description>
        <paragraph>
          Number of packets observed by a Selector.
        </paragraph>
      </description>
      <units>packets</units>
    </field>

    <field name="packetsSelected" dataType="unsigned64"
           fieldId="319" status="current" group="common">
      <description>
        <paragraph>
          Number of packets selected by a Selector.
        </paragraph>
      </description>
      <units>packets</units>
    </field>

    <field name="fixedError" dataType="float64"
           fieldId="320" status="current" group="common">
      <description>
        <paragraph>
           Specifies the maximum possible positive or negative error
           interval of the reported value for a given Information
           Element.
        </paragraph>
      </description>
      <units>
        The units of the Information Element
        for which the error is specified.
      </units>
    </field>

    <field name="relativeError" dataType="float64"
```

```
           fieldId="321" status="current" group="common">
      <description>
        <paragraph>
          Specifies the maximum possible positive or negative
          error ratio for a given Information Element.
        </paragraph>
      </description>
    </field>

    <field name="observationTimeSeconds"
           dataType="dateTimeSeconds"
           fieldId="322" status="current" group="common">
      <description>
        <paragraph>
          The absolute time of an observation.
        </paragraph>
      </description>
      <units>seconds</units>
    </field>

    <field name="observationTimeMilliseconds"
           dataType="dateTimeMilliseconds"
           fieldId="323" status="current" group="common">
      <description>
        <paragraph>
          The absolute time of an observation.
        </paragraph>
      </description>
      <units>milliseconds</units>
    </field>

    <field name="observationTimeMicroseconds"
           dataType="dateTimeMicroseconds"
           fieldId="324" status="current" group="common">
      <description>
        <paragraph>
          The absolute time of an observation.
        </paragraph>
      </description>
      <units>microseconds</units>
    </field>

    <field name="observationTimeNanoseconds"
           dataType="dateTimeNanoseconds"
           fieldId="325" status="current" group="common">
      <description>
        <paragraph>
          The absolute time of an observation.
```

```
        </paragraph>
      </description>
      <units>nanoseconds</units>
    </field>

    <field name="digestHashValue" dataType="unsigned64"
           fieldId="326" status="current" group="common">
      <description>
        <paragraph>
          The value from the digest hash function.
        </paragraph>
      </description>
    </field>

    <field name="hashIPPayloadOffset" dataType="unsigned64"
           fieldId="327" status="current" group="common">
      <description>
        <paragraph>
          The IP payload offset used by a hash based Selector.
        </paragraph>
      </description>
    </field>

    <field name="hashIPPayloadSize" dataType="unsigned64"
           fieldId="328" status="current" group="common">
      <description>
        <paragraph>
          The IP payload size used by a hash based Selector.
        </paragraph>
      </description>
    </field>

    <field name="hashOutputRangeMin" dataType="unsigned64"
           fieldId="329" status="current" group="common">
      <description>
        <paragraph>
          A value for the beginning of a hash function's
          potential output range.
        </paragraph>
      </description>
    </field>

    <field name="hashOutputRangeMax" dataType="unsigned64"
           fieldId="330" status="current" group="common">
      <description>
        <paragraph>
          A value for the end of a hash function's
          potential output range.
```

```
        </paragraph>
      </description>
    </field>

    <field name="hashSelectedRangeMin" dataType="unsigned64"
           fieldId="331" status="current" group="common">
      <description>
        <paragraph>
          A value for the beginning of a hash function's
          selected range.
        </paragraph>
      </description>
    </field>

    <field name="hashSelectedRangeMax" dataType="unsigned64"
           fieldId="332" status="current" group="common">
      <description>
        <paragraph>
          A value for the end of a hash function's
          selected range.
        </paragraph>
      </description>
    </field>

    <field name="hashDigestOutput" dataType="boolean"
           fieldId="333" status="current" group="common">
      <description>
        <paragraph>
          A boolean value, TRUE if the output from this hash Selector
          has been configured to be included in the packet report as a
          packet digest, else FALSE.
        </paragraph>
      </description>
    </field>

    <field name="hashInitialiserValue" dataType="unsigned64"
           fieldId="334" status="current" group="common">
      <description>
        <paragraph>
          The initialiser value to the hash function.
        </paragraph>
      </description>
    </field>

  </fieldDefinitions>
```

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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