

Internet-Draft  
[draft-ietf-ipfix-reducing-redundancy-00.txt](#)  
Expires: March 3, 2007

E. Boschi  
Hitachi Europe  
L. Mark  
Fraunhofer FOKUS  
B. Claise  
Cisco Systems

August 30, 2006

**Reducing redundancy in IPFIX and PSAMP reports**  
**[draft-ietf-ipfix-reducing-redundancy-00.txt](#)**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on March 3, 2007.

Copyright Notice

Copyright (C) The Internet Society (2006).

# Reducing Redundancy in IPFIX and PSAMP reports

## Abstract

This document describes a bandwidth saving method for exporting flow or packet information using the IP Flow Information Export (IPFIX) protocol. As the PSAMP protocol is based on IPFIX, these considerations are valid for PSAMP exports as well.

This method works by separating information common to several flow records from information specific to an individual flow record. Common flow information is exported only once in a data record defined by an option template, while the rest of the specific flow information is associated with the common information via a unique identifier.

## 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

<a href="#">1.</a>	<a href="#">Introduction.....</a>	<a href="#">4</a>
<a href="#">1.1</a>	<a href="#">IPFIX Documents Overview.....</a>	<a href="#">4</a>
<a href="#">1.2</a>	<a href="#">PSAMP Documents Overview.....</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Terminology.....</a>	<a href="#">5</a>
<a href="#">2.1</a>	<a href="#">Terminology Summary Table.....</a>	<a href="#">10</a>
<a href="#">2.2</a>	<a href="#">IPFIX Flows versus PSAMP Packets.....</a>	<a href="#">11</a>
<a href="#">3.</a>	<a href="#">Problem Statement and High Level Solution.....</a>	<a href="#">11</a>
<a href="#">3.1</a>	<a href="#">Per Flow Data Reduction.....</a>	<a href="#">11</a>
<a href="#">3.1.1</a>	<a href="#">Unique Data Reduction.....</a>	<a href="#">11</a>
<a href="#">3.1.2</a>	<a href="#">Multiple Data Reduction.....</a>	<a href="#">12</a>
<a href="#">3.2</a>	<a href="#">Per Packet Data Reduction.....</a>	<a href="#">14</a>
<a href="#">4.</a>	<a href="#">Specifications for bandwidth saving information export..</a>	<a href="#">16</a>
<a href="#">4.1</a>	<a href="#">Per Flow Data Reduction.....</a>	<a href="#">16</a>
<a href="#">4.1.1</a>	<a href="#">Unique Data Reduction.....</a>	<a href="#">16</a>
<a href="#">4.1.2</a>	<a href="#">Multiple Data Reduction.....</a>	<a href="#">18</a>
<a href="#">4.2</a>	<a href="#">Per-Packet Data Reduction.....</a>	<a href="#">19</a>
<a href="#">5.</a>	<a href="#">Transport Protocol Choice.....</a>	<a href="#">19</a>
<a href="#">5.1</a>	<a href="#">SCTP.....</a>	<a href="#">19</a>
<a href="#">5.2</a>	<a href="#">UDP.....</a>	<a href="#">20</a>
<a href="#">5.3</a>	<a href="#">TCP.....</a>	<a href="#">20</a>
<a href="#">6.</a>	<a href="#">commonPropertiesID Management.....</a>	<a href="#">20</a>
<a href="#">7.</a>	<a href="#">The Collecting Process Side.....</a>	<a href="#">21</a>
<a href="#">7.1</a>	<a href="#">SCTP.....</a>	<a href="#">22</a>
<a href="#">7.2</a>	<a href="#">UDP.....</a>	<a href="#">22</a>

Boschi, Mark, Claise

Expires March 2007

[Page 2]

## Reducing redundancy in IPFIX and PSAMP reports

<a href="#">7.3</a>	<a href="#">TCP</a>	<a href="#">22</a>
<a href="#">8</a>	<a href="#">Export and Evaluation Considerations</a>	<a href="#">22</a>
<a href="#">8.1</a>	<a href="#">Transport Protocol Choice</a>	<a href="#">23</a>
<a href="#">8.2</a>	<a href="#">Reduced Size Encoding</a>	<a href="#">23</a>
<a href="#">8.3</a>	<a href="#">CommonPropertiesID vs. TemplateID scope</a>	<a href="#">23</a>
<a href="#">8.4</a>	<a href="#">Efficiency Gain</a>	<a href="#">23</a>
<a href="#">9</a>	<a href="#">IANA Considerations</a>	<a href="#">24</a>
<a href="#">10</a>	<a href="#">Security Considerations</a>	<a href="#">24</a>
<a href="#">11</a>	<a href="#">Appendix A: Examples</a>	<a href="#">24</a>
<a href="#">11.1</a>	<a href="#">Per Flow Data Reduction</a>	<a href="#">24</a>
<a href="#">11.1.1</a>	<a href="#">Unique Data Reduction</a>	<a href="#">24</a>
<a href="#">11.1.2</a>	<a href="#">Multiple Data Reduction</a>	<a href="#">28</a>
<a href="#">11.2</a>	<a href="#">Per-Packet Information Export</a>	<a href="#">31</a>
<a href="#">12</a>	<a href="#">References</a>	<a href="#">34</a>
<a href="#">12.1</a>	<a href="#">Normative References</a>	<a href="#">34</a>
<a href="#">12.2</a>	<a href="#">Informative References</a>	<a href="#">35</a>
<a href="#">13</a>	<a href="#">Author's Addresses</a>	<a href="#">35</a>
<a href="#">14</a>	<a href="#">Intellectual Property Statement</a>	<a href="#">36</a>
<a href="#">15</a>	<a href="#">Copyright Statement</a>	<a href="#">36</a>
<a href="#">16</a>	<a href="#">Disclaimer</a>	<a href="#">37</a>



## **1. Introduction**

The IPFIX working group has specified a protocol to export IP Flow information [[IPFIX-PROTO](#)]. This protocol is designed to export information about IP traffic flows and related measurement data, where a flow is defined by a set of key attributes (e.g. source and destination IP address, source and destination port, etc.). However, thanks to its template mechanism, the IPFIX protocol can export any type of information, as long as the information element is specified in [[IPFIX-INFO](#)] or registered with IANA.

Regardless of the flow attributes content, flow records with common attributes export the same values in every single flow record. These common attributes may represent values common to a collection of flows or packets, or values that are invariant over time. The reduction of redundant data from the export stream can result in a significant reduction of the transferred data.

This draft specifies a way to export these invariant or common attributes only once, while the rest of the flow specific attributes are exported in regular data records. Unique common properties identifiers are used to link data records and the common attributes.

The proposed method is applicable to IPFIX flow and to PSAMP per packet information, without any changes to both the IPFIX and PSAMP protocol specifications.

### **1.1 IPFIX Documents Overview**

The IPFIX protocol [[IPFIX-PROTO](#)] provides network administrators with access to IP flow information. The architecture for the export of measured IP flow information out of an IPFIX exporting process to a collecting process is defined in [[IPFIX-ARCH](#)], per the requirements defined in [[RFC3917](#)]. This document specifies how IPFIX data record and templates are carried via a congestion-aware transport protocol from IPFIX exporting processes to IPFIX collecting process. IPFIX has a formal description of IPFIX information elements, their name, type and additional semantic information, as specified in [[IPFIX-INFO](#)]. Finally [[IPFIX-AS](#)] describes what type of applications can use the IPFIX protocol and how they can use the information provided. It furthermore shows how the IPFIX framework relates to other architectures and frameworks.

### **1.2 PSAMP Documents Overview**

The document "A Framework for Packet Selection and Reporting" [PSAMP-FMWK], describes the PSAMP framework for network elements to select

Boschi, Mark, Claise

Expires March 2007

[Page 4]

## Reducing redundancy in IPFIX and PSAMP reports

subsets of packets by statistical and other methods, and to export a stream of reports on the selected packets to a collector. The set of packet selection techniques (sampling, filtering, and hashing) supported by PSAMP are described in "Sampling and Filtering Techniques for IP Packet Selection" [[PSAMP-TECH](#)]. The PSAMP protocol [[PSAMP-PROTO](#)] specifies the export of packet information from a PSAMP exporting process to a PSAMP collecting process. Like IPFIX, PSAMP has a formal description of its information elements, their name, type and additional semantic information. The PSAMP information model is defined in [[PSAMP-INFO](#)]. Finally [[PSAMP-MIB](#)] describes the PSAMP Management Information Base.

## 2. Terminology

The terms in this section are in line with the IPFIX terminology section [[IPFIX-PROTO](#)], and [[PSAMP-PROTO](#)]. Note that this document selected the IPFIX definition of the term Exporting Process [[IPFIX-PROTO](#)], as this definition is more generic than the PSAMP definition [[PSAMP-PROTO](#)].

### 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. 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. In the IPFIX Message it generates, the Observation Domain includes its Observation Domain ID, which is unique per Exporting Process. That way, the Collecting Process can identify the specific Observation Domain from the Exporter that sends the IPFIX Messages. Every Observation Point is associated with an Observation Domain. It is RECOMMENDED that Observation Domain IDs are also unique per IPFIX Device.



## Reducing redundancy in IPFIX and PSAMP reports

### 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 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).

Boschi, Mark, Claise

Expires March 2007

[Page 6]

## Reducing redundancy in IPFIX and PSAMP reports

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.

### 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.

### Exporter

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

### 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. The Collecting Process might process or store received Flow Records, but such actions are out of scope for this document.

### Template

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.

### Template Record

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

Boschi, Mark, Claise

Expires March 2007

[Page 7]

## Reducing redundancy in IPFIX and PSAMP reports

### 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.



## Reducing redundancy in IPFIX and PSAMP reports

### Observed Packet Stream

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

### 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.

### Selection Process

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

### 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.

### 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.

### 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.



## Reducing redundancy in IPFIX and PSAMP reports

### CommonPropertiesID

An identifier of a set of common properties that is locally unique to an Exporting Process and to Observation Domain. This ID can be used to link to information reported in separate records. See [[IPFIX-INFO](#)] for the Information Element definition.

### Common Properties

Common Properties are a collection of one or more attributes shared by a set of different Flow Records. Each set of Common Properties is uniquely identifiable by means of a commonPropertiesID.

### Specific Properties

Specific Properties are a collection of one or more attributes reported in a Flow Record that are not included in the Common Properties defined for that Flow Record.

### [2.1](#) Terminology Summary Table.

+-----+-----+-----+		Contents	
		+-----+-----+-----+	
Set		Template	Record
+-----+-----+-----+		+-----+-----+-----+	
Data Set		/	Data Record(s)
+-----+-----+-----+		+-----+-----+-----+	
Template Set		Template Record(s)	/
+-----+-----+-----+		+-----+-----+-----+	
Options Template Set		Options Template Record(s)	/
+-----+-----+-----+		+-----+-----+-----+	

Figure 1: Terminology Summary Table

A Data Set is composed of Data Record(s). No Template Record is included. A Template Record or an Options Template Record defines the Data Record.

A Template Set contains only Template Record(s).

An Options Template Set contains only Options Template Record(s).

Boschi, Mark, Claise

Expires March 2007

[Page 10]

## **2.2 IPFIX Flows versus PSAMP Packets**

As described in [[PSAMP-PROTO](#)], 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.

For this document clarity, the term Flow Record represents a generic term expressing an IPFIX Flow Record or a PSAMP packet record, as foreseen by its definition. However, when appropriate, a clear distinction between Flow Record or packet Record will be made.

## **3. Problem Statement and High Level Solution**

Several Flow Records often share a set of common properties. Repeating the information about these common properties for every Flow Record introduces a huge amount of redundancy. This draft proposes a method to reduce this redundancy. The next section describes the generic concept. [Section 3.1.2](#) identifies that the proposed solution can be applied multiple times. [Section 3.2](#) utilizes the concept to export per-packet information.

### **3.1 Per Flow Data Reduction**

#### **3.1.1 Unique Data Reduction**

Consider a set of properties "A", e.g. common sourceAddressA and sourcePortA, equivalent for each Flow Records exported. Figure 2 shows how this information is repeated with classical IPFIX Flow Records, expressing the waste of bandwidth to export redundant information.

```
+-----+-----+-----+
| sourceAddressA | sourcePortA | <flow1 information> |
+-----+-----+-----+
| sourceAddressA | sourcePortA | <flow2 information> |
+-----+-----+-----+
| sourceAddressA | sourcePortA | <flow3 information> |
+-----+-----+-----+
| sourceAddressA | sourcePortA | <flow4 information> |
+-----+-----+-----+
|      ...      |      ...      |      ...      |
+-----+-----+-----+
```

Figure 2: Common and Specific Properties exported in the same record

Boschi, Mark, Claise

Expires March 2007

[Page 11]

## Reducing redundancy in IPFIX and PSAMP reports

Figure 3 shows how this information is exported when applying the specifications of this document. The Common Properties are separated from the Specific Properties for each Flow Record. The Common Properties would be exported only once in a specific Data Record (defined by an Option Template), while each Flow Record contains a pointer to the Common Properties A, along with its Flow specific information. In order to maintain the relationship between these sets of properties, we introduce indices (index A) for the Common Properties that are unique for all Common Properties entries within an Observation Domain. The purpose of the indices is to serve as a "key" identifying "rows" of the Common Properties table. The rows are then referenced by the Specific Properties by using the appropriate value for the Common Properties identifier.

index for properties A	sourceAddressA	sourcePortA
...	...	...

index for properties A	<flow1 information>
index for properties A	<flow2 information>
index for properties A	<flow3 information>
index for properties A	<flow4 information>

Figure 3: Common and Specific Properties exported in different records

This unique export of the Common Properties results in a decrease of the bandwidth requirements from the Exporter to the Collector.

### 3.1.2 Multiple Data Reduction

A Flow Record can refer to one or more Common Properties sets; the use of multiple Common Properties can lead to more efficient exports. Note that in the case of multiple Common Properties, the different sets of Common Properties MUST be disjoint (i.e. MUST not have information elements in common), to avoid potential collisions.

Boschi, Mark, Claise

Expires March 2007

[Page 12]

## Reducing redundancy in IPFIX and PSAMP reports

Consider a set of properties "A", e.g. common sourceAddressA and sourcePortA and another set of properties "B", e.g. destinationAddressB and destinationPortB. Figure 4 shows how this information is repeated with classical IPFIX export in several Flow Records.

```
+-----+-----+-----+-----+-----+
|srcAddrA|srcPortA|destAddrB|destPortB| <flow1 information> |
+-----+-----+-----+-----+-----+
|srcAddrA|srcPortA|destAddrB|destPortB| <flow2 information> |
+-----+-----+-----+-----+-----+
|srcAddrA|srcPortA|destAddrB|destPortB| <flow3 information> |
+-----+-----+-----+-----+-----+
|srcAddrA|srcPortA|destAddrB|destPortB| <flow4 information> |
+-----+-----+-----+-----+-----+
|  ...  |  ...  |  ...  |  ...  |  ...  |
+-----+-----+-----+-----+-----+
```

Figure 4: Common and Specific Properties exported in the same record

We can separate the Common Properties into the properties A composed of sourceAddressA and sourcePortA, and into the properties B composed of destinationAddressB and destinationPortB. The Flow Record that only contain the property A will only contain the index for property A, the Flow Record that only contain the property B will contain the index for property B, while the Flow Record that contain both the properties A and B contains both indexes (see Figure 5).

```
+-----+-----+-----+
| index for prop. A | sourceAddressA | sourcePortA |
+-----+-----+-----+

+-----+-----+-----+
| index for prop. B | destinationAddressB | destinationPortB |
+-----+-----+-----+

+-----+-----+-----+
|index for prop. A|index for prop. B| <flow1 information> |
+-----+-----+-----+
|index for prop. A|index for prop. B| <flow2 information> |
+-----+-----+-----+
|index for prop. A|index for prop. B| <flow3 information> |
+-----+-----+-----+
|index for prop. A|index for prop. B| <flow4 information> |
+-----+-----+-----+
|  ...  |  ...  |  ...  |
+-----+-----+-----+
```

Boschi, Mark, Claise

Expires March 2007

[Page 13]

Figure 5: Multiple Common (above) and Specific Properties (below) exported in different records

The advantage of the multiple Common Properties is that the objective of reducing the bandwidth is met while the number of index is kept to a minimum. Indeed, an alternative solution would have been to have an extra index for the property C, composed of sourceAddressA, sourcePortA, destinationAddressB, destinationPortB.

### **3.2 Per Packet Data Reduction**

The PSAMP protocol can be used for the export of per-packet information. In this case the specific packet of observation could be considered a special case of a Flow (a Flow Record composed of a single packet) and consequently per-packet information could be exported using Flow Records. However, if filtering is applied to select a subset of all packets, using IPFIX to export per-packet information is relatively inefficient since all packets belonging to the same series share common attributes (e.g. source address, destination address, etc).

A first example of the per packet data reduction is the measurement of One-Way Delay (OWD), where the exact same specific packet must be observed at the source and destination of the path to be measured. By subtracting the time of observation of the same packet at the two end points with synchronized clocks, the OWD is computed. As the OWD is measured for a specific application on which a Service Level Agreement (SLA) is bound, this translates into the observation of packets with specific properties, results of filtering. For example, all the packets of a specific source and destination IP addresses, of a specific DSCP value, and of a specific destination transport port. In order to match the identical packet at both Observation Points, a series of packets with those properties must be observed on both ends of the measurements. This implies the export of a series of Flow Records composed of two types of information: some common information for all packets, and some unique information about packets in order to generate a unique identifier for each packet passing this Observation Point (for example, a hash value on the invariant fields of the packet). So, the two IPFIX Devices composing the measurements end points can individually apply the redundancy technique described in this draft in order to save some bandwidth for the Flow Records export.

A second example of per packet data reduction is trajectory sampling. In trajectory sampling a packet is sampled on every link traversed in a measurement domain.

Boschi, Mark, Claise

Expires March 2007

[Page 14]

## Reducing redundancy in IPFIX and PSAMP reports

[TODO: make the distinction between 1. temporal export of same information from one PSAMP device 2. export of similar information from different devices. The method in this document only applies to 1.]

A third example of per packet data reduction is One-packet flows exported from a single router with a zero second export.

[TODO: This would be an example of the I.E. 313, ipHeaderPacketSection and I.E 314,ipPayloadPacketSection in PSAMP]

Figure 6, which displays the high level solution for the per-packet reduction, depicts three packets belonging to Flow A (and therefore sharing the set of Common Properties A) and one packet belonging to Flow B, respectively. It shows export records containing packet specific information and the Common Properties (source and destination address). The Common Properties introduce a huge amount of redundancy, as they are repeated for every packet in every Data Record.

```
+-----+-----+-----+
| srcAddrA | destAddrA | <packet1 information> |
+-----+-----+-----+
| srcAddrA | destAddrA | <packet2 information> |
+-----+-----+-----+
| srcAddrB | destAddrB | <packet3 information> |
+-----+-----+-----+
| srcAddrA | destAddrA | <packet4 information> |
+-----+-----+-----+
```

Figure 6: Common and Specific Properties represented in one record

In Figure 7 we separate Common Properties from Specific Properties, i.e. Common Properties from specific packet information. In order to maintain the relation between Specific (Packet) Properties and Common Properties we introduce indices (index A and index B), as previously explained.

```
+-----+-----+-----+
| srcAddrA | destAddrA | index for properties A |
+-----+-----+-----+
| srcAddrB | destAddrB | index for properties B |
+-----+-----+-----+
```



## Reducing redundancy in IPFIX and PSAMP reports

```
+-----+-----+
| index for properties A | <packet1 info> |
+-----+-----+
| index for properties A | <packet2 info> |
+-----+-----+
| index for properties B | <packet3 info> |
+-----+-----+
| index for properties A | <packet4 info> |
+-----+-----+
```

Figure 7: Common and Specific (packet) Properties exported separately

### 4. Specifications for bandwidth saving information export

The IPFIX protocol [[IPFIX-PROTO](#)] is Template based. Templates define how data should be exported, describing data fields together with their type and meaning. IPFIX specifies two types of Templates: the Template Record and the Options Template Record. The difference between the two is that the Options Template Record includes the notion of scope, defining how to scope the applicability of the Data Record. The scope, which is only available in the Options Template Record, gives the context of the reported Information Elements in the Data Records. The Template Records and Options Template Records are necessary to decode the Data Records. Indeed, by only looking at the Data Records themselves, this is impossible to distinguish a Data Record defined by Template Record from a Data Record defined by an Option Template Record. To export information more efficiently, this specification proposes to group Flow Records by their common properties. We define Common Properties as a collection of attributes shared by a set of different Flow Records.

#### 4.1 Per Flow Data Reduction

##### 4.1.1 Unique Data Reduction

As explained in Figure 8, the information is split into two parts, using two different Data Records. Common Properties MUST be exported via Data Records defined by an Option Template Record and MUST be sent only once with SCTP and TCP. These properties represent values common to several Flow Records (e.g. IP source and destination address). The Common Properties Data Records MUST be sent prior to the corresponding Specific Properties Data Records.

Boschi, Mark, Claise

Expires March 2007

[Page 16]

## Reducing redundancy in IPFIX and PSAMP reports

The Data Records reporting Specific Properties MUST be associated with the Data Records reporting the Common Properties using a unique identifier for the Common Properties, the `commonPropertiesID` Information Element. The `commonPropertiesID` MUST be exported as the scope in the Options Template Record, and also exported in the associated Template Record.

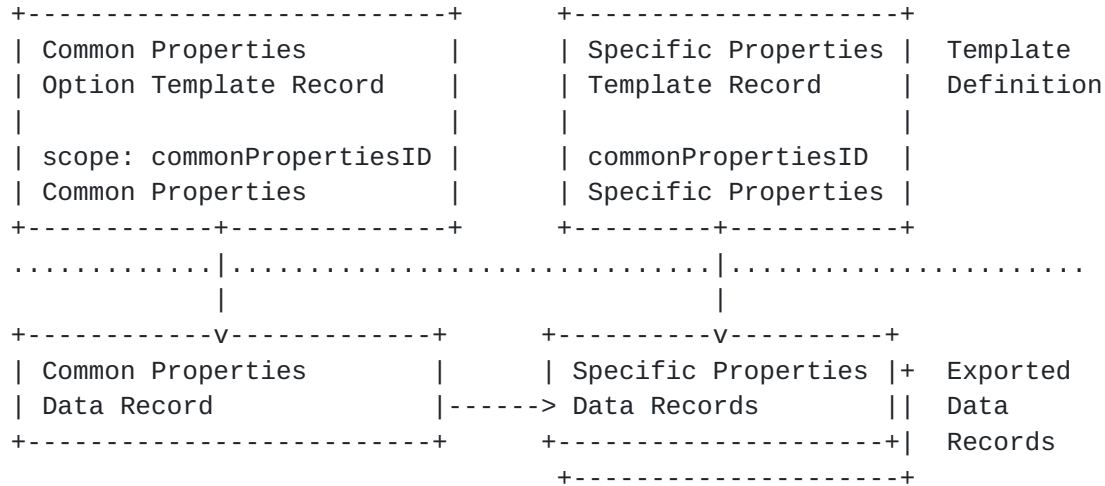


Figure 8: Template Record and Data Record dependencies

The Common Properties are valid for all Flow Records containing the associated `commonPropertiesID`. Since the `commonPropertiesID` is a 64-bit data type, this method limits the number of active data reduction to  $2^{64}$  per Exporting Process and Observation Domain. The assignment of Flow Records to common attributes could be alternatively provided by the `templateID` Information Element (instead of the `commonPropertiesID` Information Element). In this case, the scope in the Common Properties Option Template Record must contain the Template ID used in the Specific Properties Template Record, as displayed in Figure 9. The Common Properties are valid for all data records of the specified Template. In this case the use of `commonPropertiesID` is not required.



```

+-----+-----+
| Common Properties | | Specific Properties | Template
| Option Template Record | | Template Record | Definition
| | | |
| scope: Template ID | | Specific Properties |
| Common Properties | | |
+-----+-----+
.....|.....|.....
|
+-----v-----+ +-----v-----+
| Common Properties | | Specific Properties | + Exported
| Data Record |-----> Data Records || Data
+-----+-----+ +-----+-----+ | Records
+-----+-----+

```

### 4.1.2 Multiple Data Reduction

```

+-----+
| Common Properties | | Specific Properties | Template
| Option Template Record | | Template Record | Definition
| | | |
| Scope: commonPropertiesID1 | | commonPropertiesID1 |
| Scope: commonPropertiesID2 | | commonPropertiesID2 |
| Common Properties | | Specific Properties |
+-----+
.....|.....
|
+-----v-----+ +-----v-----+
| Common Properties | | Specific Properties | + Exported
| Data Record |-----> Data Records || Data
+-----+ +-----+ | Records
+-----+

```

Figure 10: Multiple data reduction



## 4.2 Per-Packet Data Reduction

From the IPFIX protocol, there are no differences between the Flow Record or per packet record data reduction, except maybe the terminology where the Specific Properties could be called packet specific properties in the following Figure 11.

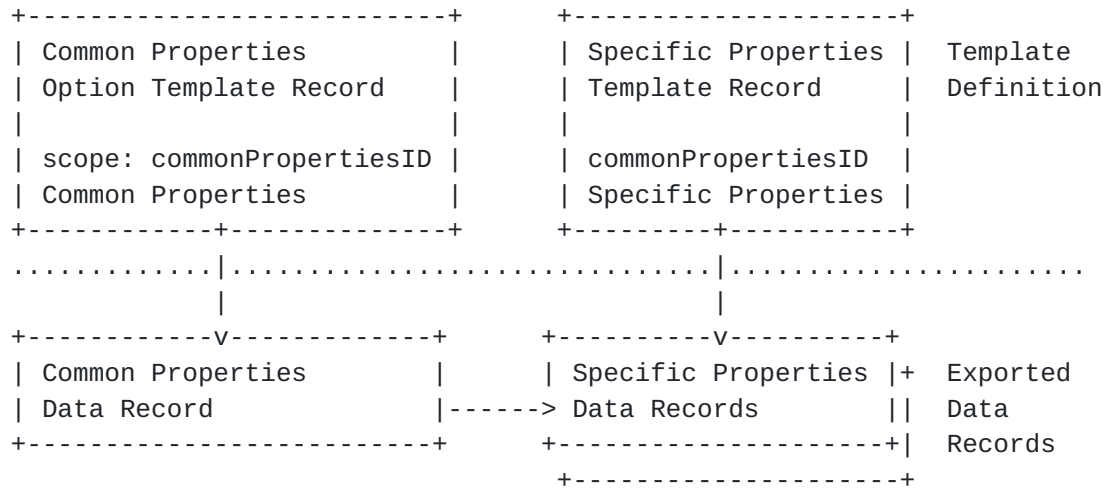


Figure 11: Per-packet data reduction

## 5. Transport Protocol Choice

This document follows the IPFIX transport protocol specifications defined in [[IPFIX-PROTO](#)]. However, depending on the transport protocol choice, this document imposes some more constraints. If SCTP is selected as the IPFIX protocol, the SCTP sub-section specifications MUST be respected. If UDP is selected as the IPFIX protocol, the UDP sub-section specifications MUST be respected. If TCP is selected as the IPFIX protocol, the TCP sub-section specifications MUST be respected.

### 5.1 SCTP

The active Common Properties MUST be sent after the SCTP association establishment before the corresponding Specific Properties Data Records. In case of SCTP association re-establishment, all active Common Properties MUST be re-sent before the corresponding Specific Properties Data Records.

The Common Properties Flow Records MUST be sent on a reliable SCTP stream.



## **5.2 UDP**

Common Properties Data Records MUST be re-sent at regular intervals, whose frequency MUST be configurable.

CommonPropertiesIDs have a specified lifetime during which they cannot be reused. After that time a commonPropertiesID can be assigned to another set of Common Properties. CommonPropertiesID whose lifetime has longer expired SHOULD be preferred. The lifetime MUST be configurable.

## **5.3 TCP**

Common Properties MUST be sent after the TCP connection establishment before the corresponding Specific Properties Data Records. In case of TCP connection re-establishment, all active Common Properties MUST be re-sent before the corresponding Specific Properties Data Records.

## **6. commonPropertiesID Management**

The commonPropertiesID is an identifier of a set of common properties that is locally unique to an Exporting Process and to Observation Domain. The Exporting Process MUST manage the commonPropertiesIDs allocations for its Observation Domains. Different Observation Domains from the same Exporter MAY use the same commonPropertiesID value to refer to different sets of Common Properties.

The commonPropertiesID values MAY be assigned sequentially, but it's NOT REQUIRED. Particular commonPropertiesID ranges or values MAY have explicit meanings for the IPFIX Device. For example, commonPropertiesID values may be assigned based on the result of a hash function, etc...

Using a 64-bit commonPropertiesID Information Element allows the export of  $2^{64} - 1$  active sets of Common Properties, per Observation Domain, per Exporting Process.

CommonPropertiesIDs that are not used anymore SHOULD be withdrawn. The Common Properties ID withdrawal message is an Option Data Record consisting of only one scope field namely the CommonPropertiesID and no non-scope fields.



## Reducing redundancy in IPFIX and PSAMP reports

```
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 = 14 octets          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Template ID = 259    |          Field Count = 1          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Scope Field count = 1  |0|  commonPropertiesID = XX      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|  Scope 1 Field Length = 8      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Figure 12: CommonPropertiesID withdrawal template

```
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 = 259          |          Length = 12 octets          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                N                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                ...                               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Figure 13: CommonPropertiesID withdrawal record, withdrawing CommonPropertiesID N

## 7. The Collecting Process Side

The Collecting Process can either store the Flow Records as they arrive, without reconstructing the initial Flow Record, or reconstruct the initial Flow Record. In the former case there might be less storage capacity required at the Collector side. In the latter the collector job is more complex and time-consuming due to the higher resource demand for record processing in real time.

Like TemplateIDs the CommonPropertiesIDs are generated dynamically by the Exporting Process. The CommonPropertiesIDs are only valid within the protocol stack. Hence a restart of the exporting process may imply a renumbering of CommonPropertiesIDs. For this reason it is not recommended to use the CommonPropertiesIDs outside the protocol stack e.g. to store them within a database. Outside the protocol stack there is additional information needed to keep a non-ambiguous association between the related Common Properties and Specific Properties.



## Reducing redundancy in IPFIX and PSAMP reports

If the Collecting Process has received the Specific Properties Data Record before the associated Common Properties Data Record, the Collecting Process MAY store the Specific Properties Data Record and await the retransmission or out-of-order arrival of the Common Properties Data Record.

If a Collection Process receives a CommonPropertiesID Withdraw Record, the Collection Process MUST expire the related Common Properties data.

If SCTP is selected as the IPFIX protocol, the SCTP sub-section specifications MUST be respected. If UDP is selected as the IPFIX protocol, the UDP sub-section specifications MUST be respected. If TCP is selected as the IPFIX protocol, the TCP sub-section specifications MUST be respected.

### [7.1](#) SCTP

When the SCTP association is reset, either gracefully or abnormally, the Collecting Processes MUST delete all commonPropertiesID values associated with that association.

### [7.2](#) UDP

The Collecting Process associates a lifetime with each commonPropertiesID. The mapping of Data Records to Common Properties uses the most recent Common Properties definition associated to the specified commonPropertiesID. The lifetime of the CommonPropertiesID ends on the receipt of a CommonPropertiesID withdrawal record. If there is no flow definition associated with that commonPropertiesID or the lifetime of the flow definition has expired, no mapping is possible. In this case the Collecting Process MAY store the Specific Properties and await the retransmission or out-of-order arrival of the Common Properties.

### [7.3](#) TCP

When the TCP connection is reset, either gracefully or abnormally, the Collecting Processes MUST expire all commonPropertiesID values corresponding to that connection.

## [8. Export and Evaluation Considerations](#)

The main advantage of the method specified in this document is the reduction in the amount of measurement data that has to be transferred from the Exporter to the Collector. In addition there might be less storage capacity required at the Collector side if the



## Reducing redundancy in IPFIX and PSAMP reports

Collector decides to store the Flow Records as they arrive, without reconstructing the initial Flow Record.

On the other hand, these methods require additional resources on both the Exporter and the Collector. The Exporter has to manage Common Properties information and to assign `commonPropertiesId` values to Flow Records. The Collector has to process records described by two templates instead of just one. Additional effort is also required when post processing the measurement data, in order to correlate Flow Records with Common Properties information.

### **8.1 Transport Protocol Choice**

The proposed method is most effective using a reliable transport protocol for the transfer of the Common Properties. Therefore the use of SCTP or TCP is recommended. However, if the path from the Exporting Process to the Collecting Process is not fully reliable, the SCTP or TCP retransmission might reduce the benefits of this specification. If the path from the Exporting Process to the Collecting Process is full reliable, the use of UDP is less effective because the common properties have to be re-sent regularly.

### **8.2 Reduced Size Encoding**

The transfer of the `CommonPropertiesIDs` originates some overhead. Note that IPFIX allows reduced-size encoding of Information Elements. In cases where the range of the `commonPropertiesID` can be restricted, reduced-size encoding can be applied also to the `commonPropertiesID`, and would result in a further bandwidth efficiency gain.

### **8.3 `CommonPropertiesID` vs. `TemplateID` scope**

The assignment of Flow Records to common attributes could be done via the `CommonPropertiesID` and alternatively via the `templateID` Information Element. In the second case the `commonPropertiesID` is not required: this reduces the overhead but the Exporting Process must use one `templateID` per set of Common Properties. In the general case, this method is not scalable, but it can be suitable for certain applications.

### **8.4 Efficiency Gain**

The example in [section 11.2](#) below uses IPFIX to export measurement data for each received packet. In that case, for a flow of 1000 packets the amount of data can be decreased more than 33 percent.

Boschi, Mark, Claise

Expires March 2007

[Page 23]

## Reducing redundancy in IPFIX and PSAMP reports

While the goal of this specification is to reduce the bandwidth, the efficiency might be limited. Indeed, the efficiency gain is based on the numerous redundant information in flows. While the Exporting Process can evaluate the direct gain for the Flow Records to be exported, it can't predict whether future Flow Records would contain the information specified by active commonPropertiesID values. This implies that the efficiency factor of this specification is higher for specific applications where filtering is involved, such as one-way delay or trajectory sampling.

### 9. IANA Considerations

This document has no actions for IANA.

### 10. Security Considerations

For the proposed use of the IPFIX protocol for bandwidth-saving export the security considerations as for the IPFIX protocol apply.

### 11. [Appendix A](#): Examples

#### 11.1 Per Flow Data Reduction

##### 11.1.1 Unique Data Reduction

In this section we show how flow information can be exported efficiently using the method described in this draft. Let's suppose we have to periodically export data about two IPv6 Flows.

In this example we report the following information:

Flow	dstIPv6Address	dst- Port	nPkts	nBytes
A	5F05:2000:80AD:5800:0058:0800:2023:1D71	80	30	6000
A	5F05:2000:80AD:5800:0058:0800:2023:1D71	80	50	9500
B	5F05:2000:80AD:5800:0058:00AA:00B7:AF2B	1932	60	8000
A	5F05:2000:80AD:5800:0058:0800:2023:1D71	80	40	6500



## Reducing redundancy in IPFIX and PSAMP reports

A	5F05:2000:80AD:5800:0058:0800:2023:1D71	80	60	9500
B	5F05:2000:80AD:5800:0058:00AA:00B7:AF2B	1932	54	7600

The Common Properties in this case are the destination IPv6 address and the destination port. We first define an Option Template that contains the following Information Elements:

- Scope: the commonPropertiesID, with a type of 137 [[IPFIX-INFO](#)] and a length of 8 octets.
- The destination IPv6 address, destinationIPv6Address [IPFIX-INFO], with a type of 28 and a length of 16 octets
- The destination port, destinationTransportPort [[IPFIX-INFO](#)] with a type of 11, and a length of 2 octets

Figure 14 shows the Option template defining the Common Properties with commonPropertiesID as scope:

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 = 24 octets
+--+			
	Template ID = 257		Field Count = 3
+--+			
	Scope Field count = 1	0	commonPropertiesID = 137
+--+			
	Scope 1 Field Length = 8	0	destinationIPv6Address = 28
+--+			
	Field Length = 16	0	destinationTransportPort = 11
+--+			
	Field Length = 2		(Padding)
+--+			

Figure 14: Common Properties Option Template

The Specific Properties Template consists of the information not contained in the Option Templates, i.e. flow specific information, in this case the number of packets and the number of bytes to be reported. Additionally, this Template contains the commonPropertiesID. In Data Records, the value of this field will contain one of the unique indices of the Option Records exported



## Reducing redundancy in IPFIX and PSAMP reports

before. It contains the following Information Elements (see also Figure 15):

- commonPropertiesID with a length of 8 octets
- The number of packets of the Flow: inPacketDeltaCount in [IPFIX-INFO], with a length of 4 octets
- The number of octets of the Flow: inOctetDeltaCount in [IPFIX-INFO], with a length of 4 octets

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	2	3	4	5	6	7	8	9										
+-+-+...+-+-+ (40 dashes)																																																	
										Set ID = 2																				Length = 20 octets																			
+-+-+...+-+-+ (40 dashes)																																																	
										Template ID = 258																				Field Count = 4																			
+-+-+...+-+-+ (40 dashes)																																																	
0										commonPropertiesID = 137																				Field Length = 8																			
+-+-+...+-+-+ (40 dashes)																																																	
0										inPacketDeltaCount = 2																				Field Length = 4																			
+-+-+...+-+-+ (40 dashes)																																																	
0										inOctetDeltaCount = 1																				Field Length = 4																			
+-+-+...+-+-+ (40 dashes)																																																	

Figure 15: Specific Properties Template

Considering the data shown at the beginning of this example, the following two Data Records will be exported:

Common-PropertiesID	dstAddress	dst-Port
101	5F05:2000:80AD:5800:0058:0800:2023:1D71	80
102	5F05:2000:80AD:5800:0058:00AA:00B7:AF2B	1932

The Data Records reporting the Common Properties will look like:



Reducing redundancy in IPFIX and PSAMP reports

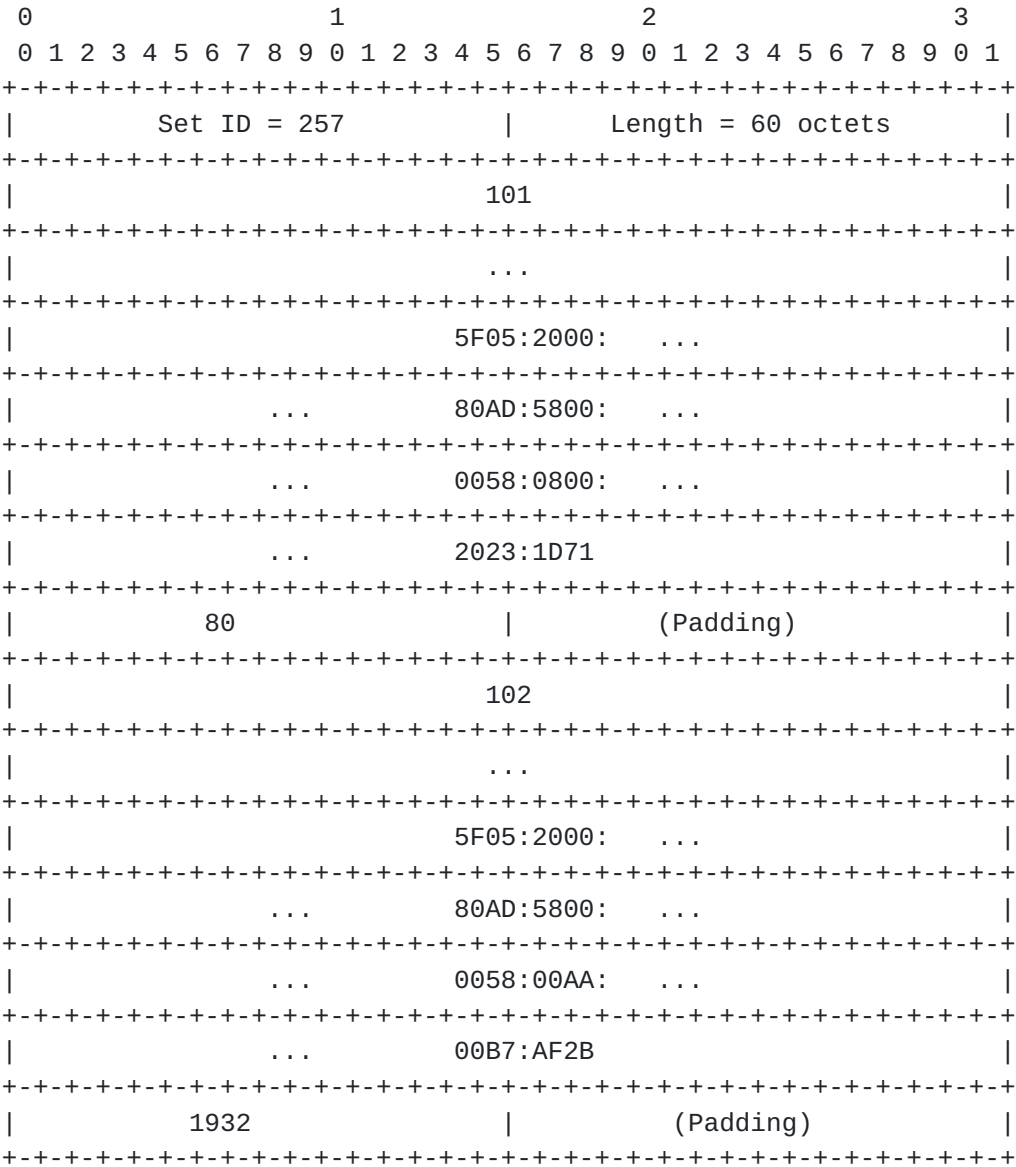


Figure 16: Data Records reporting Common Properties

The Data Records will in turn be:

commonPropertiesID	inPacketDeltaCount	inOctetDeltaCount
101	30	6000
101	50	9500
102	60	8000
101	40	6500
101	60	9500
102	54	7600

Boschi, Mark, Claise

Expires March 2007

[Page 27]

## Reducing redundancy in IPFIX and PSAMP reports

Figure 17 shows the first Data Record listed in the table:

[illegible]

Figure 17: Data Record reporting Specific Properties

### 11.1.2 Multiple Data Reduction

In this example we export the following flow information:

Flow	srcAddr	srcPort	dstAddr	dstPort	nPackets	nBytes
A	10.0.0.1	1932	10.0.1.2	80	30	6000
B	10.0.0.3	2032	10.0.1.2	80	50	9500

Figure 18 shows the Option Templates, containing the Common Properties together with the commonPropertiesID as Scope.

In the first Common Properties Option Template we export the following Information Elements:

- Scope 1: the Common Properties ID, commonPropertiesId with a type of 137 [[IPFIX-INFO](#)]. Note that the commonProperties IE has a length of 8 octets, but if smaller size is sufficient to carry any value the Exporter may need to deliver, reduced size encoding can be used. In this example we use reduced sizing, of 4 octets.
- the source IPv4 Address, sourceIPv4Address [[IPFIX-INFO](#)], with a type of 8 and a length of 4 octets
- the source Port, sourceTransportPort [[IPFIX-INFO](#)], with a type of 7 and a length of 2 octets







## Reducing redundancy in IPFIX and PSAMP reports

Considering the values given at the beginning of this section we will export the Common Properties using the following Data Records:

commonPropertiesID	sourceAddress	sourcePort
-----+-----+-----		
101	10.0.0.1	1932
102	10.0.0.3	2032

and

commonPropertiesID	dstAddress	dstPort
-----+-----+-----		
103	10.0.1.2	80

The Specific Properties Template consists of the information not contained in the Option Templates, i.e. flow specific information. Additionally, this Template contains the two commonPropertiesID. In Data Records, the values of each of these fields will contain one of the unique indices specified in the Option Records exported previously.

Figure 19 displays the Template including the commonPropertiesID plus the Specific Properties. In this example we export the following Information Elements:

- commonPropertiesID for the source fields with a length of 4 octets (reduced size encoding)
- commonPropertiesID for the destination fields with a length of 4 octets (reduced size encoding)
- the number of packets of the Flow: inPacketDeltaCount in [[IPFIX-INFO](#)], with a length of 4 octets
- the number of octets of the Flow: inOctetDeltaCount in [[IPFIX-INFO](#)], with a length of 4 octets



## Reducing redundancy in IPFIX and PSAMP reports

```

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 = 2          |          Length = 24 octets          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Template ID = 259   |          Field Count = 4            |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0| commonPropertiesID = 137   |          Field Length = 4            |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0| commonPropertiesID = 137   |          Field Length = 4            |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0|   inPacketDeltaCount = 2   |          Field Length = 4            |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|0|   inOctetDeltaCount = 1    |          Field Length = 4            |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Figure 19: Example Specific Properties Template

Considering the values given at the beginning of this section, the Data Records of the two flows will look like:

```

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 = 256        |          Length = 28                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                101                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                103                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                30                                |          6000          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                102                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                103                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                50                                |          9500          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

Figure 20: Specific Properties

### [11.2](#) Per-Packet Information Export

This section demonstrates per-packet information export to support passive One-Way Delay (OWD) measurement. The Templates required for exporting measurement data of this kind are illustrated in the figures below.







## Reducing redundancy in IPFIX and PSAMP reports

Figure 21: Example Flow Properties Template

For passive One-Way-Delay measurement, the Packet Properties Template, or Specific Properties Template, consists of at least Timestamp and Packet ID. Additionally, this template contains a commonPropertiesId field to associate the packet with a Flow.

Figure 22 displays the template with the packet properties. In this example we export the following Information Elements:

- commonPropertiesID. In this case reduced size encoding is used, and the Information Element is declared with a length of 4 octets instead of 8.
- packetTimestamp, packetID, and packetLength. Since packetTimestamp, packetID, and packetLength are not (yet) IETF-defined information elements, we export them as enterprise-specific IEs. The three IEs have respectively a type of 220, 221, and 222 and a length of 8, 4, and 4 octets.

```

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 = 2           |           Length = 36 octets       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Template ID = 257    |           Field Count = 4          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0| commonPropertiesID = 137      |           Field Length = 4         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|1| packetTimestamp = 220        |           Field Length = 8         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Enterprise number                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|1| packetID = 221               |           Field Length = 4         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Enterprise number                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|1| packetLength = 222           |           Field Length = 4         |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Enterprise number                       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Figure 22: Example Packet Properties Template



## Reducing redundancy in IPFIX and PSAMP reports

At the collection point, packet records from the two measurement points are gathered and correlated by means of the packet ID. The resulting delay data records are exported in a similar manner as the packet data. One-way delay data is associated with flow information by the commonPropertiesId field. The OWD properties contain the Packet Pair ID (which is the packet ID of the two contributing packet records), the timestamp of the packet passing the reference monitor point in order to reconstruct a time series, the calculated delay value, and the commonPropertiesID.

In this example using IPFIX to export the measurement data for each received packet 30 bytes have to be transferred (sourceAddressV4=4, destinationAddressV4=4, classOfServiceV4=1, protocolIdentifier=1, sourceTransportPort=2, destinationTransportPort=2, packetTimestamp=8, packetID=4, packetLength=4). Without considering the IPFIX protocol overhead a flow of 1000 packets produces 30000 bytes of measurement data. Using the proposed optimization each packet produces an export of only 20 bytes (packetTimestamp=8, packetID=4, packetLength=4, commonPropertiesID=4). The export of the flow information produces 18 bytes (sourceAddressV4=4, destinationAddressV4=4, classOfServiceV4=1, protocolIdentifier=1, sourceTransportPort=2, destinationTransportPort=2, commonPropertiesID=4). For a flow of 1000 packets this sums up to 20018 bytes. This is a decrease of more than 33 percent.

## **12. References**

### **12.1 Normative References**

- [RFC2119] S. Bradner, Key words for use in RFCs to Indicate Requirement Levels, [BCP 14](#), [RFC 2119](#), March 1997
- [IPFIX-PROTO] B. Claise et Al, IPFIX Protocol Specification, <[draft-ietf-ipfix-protocol-22.txt](#)>, Internet-Draft work in progress June 2006
- [IPFIX-INFO] J. Quittek, S. Bryant, B. Claise, J. Meyer, Information Model for IP Flow Information Export, <[draft-ietf-ipfix-info-12.txt](#)>, Internet-draft work in progress, June 2006
- [PSAMP-PROTO] B. Claise, J. Quittek, A. Johnson, PSAMP Protocol Specification, <[draft-ietf-psamp-protocol-06.txt](#)>, Internet-Draft work in progress, June 2006



## **12.2 Informative References**

- [IPFIX-ARCH] G. Sadasivan, N. Brownlee, B. Claise, J. Quittek, Architecture Model for IP Flow Information Export, <[draft-ietf-ipfix-arch-11.txt](#)>, Internet-Draft work in progress, June 2006
- [IPFIX-AS] T. Zseby, E. Boschi, N. Brownlee, B. Claise, IPFIX Applicability, <[draft-ietf-ipfix-as-10.txt](#)>, Internet-Draft work in progress, August 2006
- [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](#)>, Internet-Draft work in progress, January 2006
- [PSAMP-INFO] T. Dietz, F. Dressler, G. Carle, B. Claise, P. Aitken, Information Model for Packet Sampling Exports, <[draft-ietf-psamp-info-04.txt](#)>, Internet-Draft work in progress, March 2006
- [PSAMP-MIB] T. Dietz, B. Claise, Definitions of Managed Objects for Packet Sampling, <[draft-ietf-psamp-mib-06.txt](#)>, Internet-Draft work in progress, June 2006
- [PSAMP-FMWK] N. Duffield, D. Chiou, B. Claise, A. Greenberg, M. Grossglauser, P. Marimuthu, J. Rexford, G. Sadasivan, A Framework for Passive Packet Measurement, <[draft-ietf-psamp-framework-10.txt](#)>, Internet-Draft work in progress, January 2005
- [RFC3917] J. Quittek, T. Zseby, B. Claise, S. Zander, Requirements for IP Flow Information Export, [RFC 3917](#), October 2004

## **13. Author's Addresses**

Elisa Boschi  
Hitachi Europe SAS  
Immeuble Le Theleme  
1503 Route des Dolines  
06560 Valbonne, France  
Phone: +33 4 89874180  
Email: elisa.boschi@hitachi-eu.com



## Reducing redundancy in IPFIX and PSAMP reports

Lutz Mark  
Fraunhofer Institute for Open Communication Systems  
Kaiserin-Augusta-Allee 31  
10589 Berlin  
Germany  
Phone: +49-30-34 63 7306  
Fax: +49-30-34 53 8306  
Email: mark@fokus.fraunhofer.de

Benoit Claise  
Cisco Systems  
De Kleetlaan 6a b1  
Diegem 1813  
Belgium  
Phone: +32 2 704 5622  
Email: bclaise@cisco.com

### **14. Intellectual Property Statement**

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information 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](mailto:ietf-ipr@ietf.org).

### **15. Copyright Statement**

Copyright (C) The Internet Society (2006). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Boschi, Mark, Claise

Expires March 2007

[Page 36]

## **16. Disclaimer**

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.