

Internet Draft
Document: <[draft-pohl-perpktinfo-02.txt](#)>
Expires: August 2005

G. Pohl
Fraunhofer FOKUS
L. Mark
Fraunhofer FOKUS
E. Boschi
Fraunhofer FOKUS
February 2005

Use of IPFIX for Export of Per-Packet Information

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [section 3 of RFC 3667](#). 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 become aware will be disclosed, in accordance with [RFC 3668](#).

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>.

Pohl, Mark, Boschi

Expires August 2005

[Page 1]

Use of IPFIX for Export of Per-Packet Information

Abstract

This document describes the usage of the IP Flow Information Export (IPFIX) protocol for the case of exporting and processing per-packet information.

The main idea is to export two types of records per flow: the first one contains the usual flow information plus a unique flow identifier while the other one consists of the actual per-packet information plus a flow identifier that refers to the flow the specific packet belongs to -- that means the flow identifier is used to associate the packet data to its corresponding flow.

Table of Contents

1.	Introduction	2
2.	Terminology	2
3.	General Problem Statement	3
4.	Export Per-Packet Information	3
5.	Flow ID Management	5
6.	Example of Per-Packet Information Export	5
7.	IPFIX for per-packet information export and PSAMP	6
8.	Export and evaluation considerations	6
9.	IANA Consideration	7
10.	Security Considerations	7
11.	References	7
11.1	Normative References	7
11.2	Informative References	7
12.	Author's Addresses	8
13.	Intellectual Property Statement	8
14.	Copyright Statement	9
15.	Disclaimer	9

1. Introduction

In the scope of passive QoS Measurements, there is often the need to exchange and export measurement data in a finer granularity than per flows. One typical application is passive One-Way-Delay measurement; this draft takes it as example when demonstrating the need for information export on a per-packet basis.

The IPFIX protocol however, has been designed to export flow records. A possible approach to export packet records using IPFIX could be exporting flow records containing information about single packets. This method has been proposed by the PSAMP working group in [Cla204]. Exporting flow related information per-packet introduces a high degree of redundancy. This draft shows how packet information and flow information can be

efficiently exported and related using IPFIX.

2. Terminology

Collecting Process

The collecting process receives records of flow or packet information. The data is stored for later processing (by the calculation process)

Exporting Process

The exporting processes send flow and packet records to the collecting processes. The records are generated by the measurement process.

Filtering

Use of IPFIX for Export of Per-Packet Information

Filtering selects a subset of packets by applying deterministic functions on parts of the packet content like header fields or parts of the payload. A filtering process needs to process the packet (look at packet header and/or payload) in order to make the selection decision.

Measurement Device

A measurement device has access to at least one observation point. It is hosting at least one measurement process and one export process.

Metering/Measurement Process

The measurement process generates records of packet and flow information. Packets passing the observation point are captured, time stamped, filtered and classified. The measurement process calculates the packet IDs.

Passive One-Way-Delay Measurement

Abbreviated: POWD Measurement

3. General Problem Statement

In [\[Cla104\]](#) the IPFIX working group has defined a protocol to transport measurement data containing flow information.

The main purpose of the protocol is to exchange information about IP traffic flows. In this scope a flow is defined by a set of key attributes (source/destination address, source/destination port, Layer3 Protocol Type, TOS/DSCP byte, interface of the flow exporting network element). As such, a flow is a collection of packets that share a set of common attributes.

However, for a number of metrics there is a need to export per-packet data.

Undoubtedly a single packet could be considered a special case of a flow and thus, per-packet information could be exported using flow records. Doing this though would have consequences on the efficiency of the exporting procedure, as it would mean additional overhead. Packets belonging to the same flow share common attributes, i.e. source address, destination address, etc. Exporting these attributes on a per-packet basis, each time with a different packet ID, would be redundant information.

There are cases however, where it is desirable to keep flow information along with the per-packet information, that is, when analyzing packet characteristics while observing flows. This document proposes a solution that reduces the overhead caused by

the flow properties while keeping a link to flow information.

The proposed method does not need any changes to the IPFIX protocol.

4. Export Per-Packet Information

Figure 2 depicts three packets belonging to flow A and one packet that belongs to flow B, respectively. It shows export records containing packet information plus flow information (source and destination address). Undoubtedly, the flow information introduces a huge amount of redundancy, as it is repeated for every packet in every record. Minimizing the redundancy is a common problem in relational data base design and we apply here similar solutions to those proposed in that area.

Use of IPFIX for Export of Per-Packet Information

In Figure 2 we separate flow from packet information. In order to maintain the relation between Packet Properties and Flow Properties we introduce indices (idxA and idxB) for the Flow Properties that are unique for all Flow Property entries. The purpose of the indices is to serve as "primary key" that identifies rows of the Flow Properties. The rows are then referenced by the Packet Properties by using the appropriate value for the flow identifier.

One-packet flows

```
+-----+-----+-----+-----+
| srcA | dstA | packet info| ... |
+-----+-----+-----+-----+
| srcA | dstA | packet info| ... |
+-----+-----+-----+-----+
| srcB | dstB | packet info| ... |
+-----+-----+-----+-----+
| srcA | dstA | packet info| ... |
+-----+-----+-----+-----+
```

Figure 1: Flow and packet information represented in one-packet flows

	Packet Properties
Flow Properties	+-----+-----+-----+-----+
	>idxA packet info ...
+-----+-----+-----+	+-----+-----+-----+-----+
srcA dstA idxA <	>idxA packet info ...
+-----+-----+-----+	+-----+-----+-----+-----+
srcB dstB idxB <-----	>idxB packet info ...
+-----+-----+-----+	+-----+-----+-----+-----+
	>idxB packet info ...
	+-----+-----+-----+-----+

here, the linkage of one packet and
flow B (srcB, dstB, idxB) is explicitly drawn

Figure 2: Flow information and packet information

The IPFIX protocol is template based like NetFlow version 9. For a complete description of features of IPFIX refer to [Fehler! Verweisquelle konnte nicht gefunden werden.].

Templates define the structure of data to be exported, including describing data fields to be exported together with their type and meaning.

For Measurement Process Results we define two different template records, namely Flow Properties and Packet Properties. The Flow

Properties templates SHOULD be sent before the Packet Properties Templates.

In Figure 3, the Flow Properties template defines the attributes for a flow; e.g. IP source and destination address and the flow identifier. The flow identifier is a unique identifier for flow definitions; this field allows packet records to reference flow attributes. Subsequent data records of this template define the actual flows.

The format for the information related to single packets is defined in the Packet Properties template. This information is packet specific and normally not shared between many packets. Otherwise one would rather consider the information as flow related and therefore it needs not be exported in every record.

Use of IPFIX for Export of Per-Packet Information

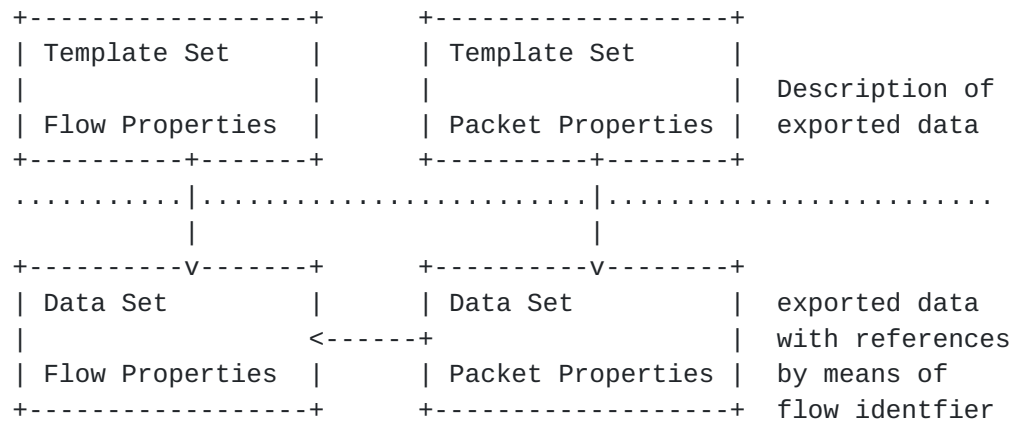


Figure 3: Template FlowSet and Data FlowSet dependencies

5. Flow ID Management

Like template IDs the flow IDs have to be unique per observation domain (source identifier in the IPFIX header). Using 32 bit flow IDs allows the export of 2^{32} active flows in parallel.

If it is desired, one can optionally use option templates to specify the mapping between two flow and packet properties templates. In this case, the flow ID only has to be unique within this association.

6. Example of Per-Packet Information Export

To demonstrate how to use IPFIX efficiently to export per-packet information, this section proposes how to use the IPFIX protocol for exporting flow information and per-packet information (in this case related to a long-lived flow) for OWD computation.

In order to acquire a One-Way path delay information, two measurement points with synchronized clocks must exist, one at each end of the path under examination. Both measurement points will capture packets, assign them timestamps and generate an identifier for a packet passing that point. Each measurement point will export its measurement data to a collecting process where the data are correlated based on the packet identifiers and timestamps and then the delay is calculated as a difference of two timestamps of a packet pair.

The templates that would be needed for exporting measurement data of this kind are illustrated in Figure 4.

The upper part of the figure shows the template containing the information concerning flows. The lower part displays the template with the packet properties.

For passive One-Way-Delay measurement, the Packet Properties template consists of at least Timestamp and Packet ID. Additionally, this template contains a flow identifier field. In packet records, the value of this field will contain one of the unique indices of the flow records exported before.

Use of IPFIX for Export of Per-Packet Information

+-----+-----+-----+..			
flowID	sourceAddressV4	destinationAddressV4	
unsigned32/vendor	ipv4Address/ID 8	ipv4Address/ID 12	
+-----+-----+-----+..			
..+-----+-----+-----+..			
classOfServiceV4	protocolIdentifier	transportSourcePort	
octet/ID 5	octet/ID 4	unsigned16/ID 7	
..+-----+-----+-----+..			
..+-----+-----+-----+..			
transportDestinationPort			
unsigned16/ID 11			
..+-----+-----+-----+..			
FlowPropTemplate			
=====			
PacketPropTemplate			
+-----+-----+-----+..			
packetTimestamp	packetID	packetLength	
unsigned64/vendor	unsigned32/vendor	unsigned32/vendor	
+-----+-----+-----+..			
..+-----+-----+-----+..			
flowID			
unsigned32/vendor			
..+-----+-----+-----+..			

Figure 4: Example Templates for Flow and Packet Properties

The delay is derived by a calculation step: At the collection point packet records of 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 have been. Especially, the linkage between delay data and flow information is represented with the discussed flow identifier fields. The OWD properties contain the Packet Pair ID (which is the packet ID matching that of the two contributing packet records), a timestamp (which is the timestamp of the packet passing the reference monitor point) in order to reconstruct a time series, the calculated delay value, and finally a flow identifier.

7. IPFIX for per-packet information export and PSAMP

In [[Cla204](#)] the PSAMP working group proposes to use IPFIX to export packet information from a PSAMP Exporting Process to a PSAMP Collecting Process. Even though no new version of the draft has been produced the solution seems to be accepted from the group.

While IPFIX is well suited for the purpose due to the good match between the IPFIX and PSAMP architectures and to the fact that IPFIX satisfies PSAMP requirements, the described approach has a high degree of redundancy. It proposes to treat packets as flows and export per-packet information using flow records.

We propose to use the solution described in this draft to efficiently export PSAMP packet information.

8. Export and evaluation considerations

Use of IPFIX for Export of Per-Packet Information

The main advantage of this proposed method to export per-packet information is the reduced amount of measurement data that has to be transferred from the exporter to the collector. In addition there is less storage capacity needed at the collector side.

On the other hand there is some extra processing power needed on the exporter side to manage flow information and to assign packets to flows. The collector has to process records of two templates instead of just one but has to read and write less data. Additional effort is needed when post processing the measurement data, because now the correlation of flow and packet information is needed.

In the above example (see Figure 4) using IPFIX to export the measurement data for each received packet 28 bytes have to be transferred (sourceAddressV4=4, destinationAddressV4=4, classOfServiceV4=1, protocolIdentifier=1, transportSourcePort=2, transportDestinationPort=2, packetTimestamp=8, packetID=4, packetLength=2). Disregarding the IPFIX protocol overhead a flow of 1000 packets produces 28000 bytes of measurement data. Using the proposed optimization each packet produces an export of only 16 bytes (packetTimestamp=8, packetID=4, packetLength=2, flowID=2). The export of the flow information produces 16 bytes (sourceAddressV4=4, destinationAddressV4=4, classOfServiceV4=1, protocolIdentifier=1, transportSourcePort=2, transportDestinationPort=2, flowID =2). For a flow of 1000 packet this sums up to 16016 bytes. This is a decrease of more than 40 percent.

9. IANA Consideration

This document does not imply any IANA action.

10. Security Considerations

For the proposed use of the IPFIX protocol for export of per-packet information the security considerations as for the IPFIX protocol apply.

11. References

11.1 Normative References

- [Cla104] Benoit Claise: IPFIX Protocol Specification, IETF draft work in progress
<[draft-ietf-ipfix-protocol-05.txt](#)>, August 2004
- [Cla204] Benoit Claise: PSAMP Protocol Specification, Internet Draft <[draft-ietf-psamp-protocol-01.txt](#)>,

February 2004

11.2 Informative References

- [DuGG03] Nick Duffield, et Al.: A Framework for Packet Selection and reporting, Internet Draft
<[draft-ietf-psamp-framework-08.txt](#)>, September 2004

- [DuGr00] Nick Duffield, Matthias Grossglauser: Trajectory Sampling for Direct Traffic Observation,
Proceedings of ACM SIGCOMM 2000, Stockholm, Sweden,
August 28 - September 1, 2000

- [QuZC04] J. Quittek, et Al.: Requirements for IP Flow Information Export, Internet Draft
<[draft-ietf-ipfix-reqs-16.txt](#)>, June 2004

Use of IPFIX for Export of Per-Packet Information

- [GrDM98] Ian D. GRAHAM, Stephen F. DONNELLY, Stele MARTIN, Jed MARTENS, John G. CLEARY: Nonintrusive and Accurate Measurement of Unidirectional Delay and Delay Variation on the Internet, INET'98, Geneva, Switzerland, July 21-24, 1998
- [TPBC03] T. Zseby, E. Boschi, R. Penno, N. Brownlee, B. Claise: IPFIX Applicability, Internet Draft <[draft-ietf-ipfix-as-02.txt](#)>, July 2004

12. Author's Addresses

Elisa Boschi
Fraunhofer Institute for Open Communication Systems
Kaiserin-Augusta-Allee 31
10589 Berlin
Germany
Phone: +49-30-34 63 7366
Fax: +49-30-34 53 8366
Email: boschi@fokus.fraunhofer.de

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

Guido Pohl
Fraunhofer Institute for Open Communication Systems
Kaiserin-Augusta-Allee 31
10589 Berlin
Germany
Phone: +49-30-34 63 7164
Fax: +49-30-34 53 8164
Email: pohl@fokus.fraunhofer.de

13. 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.

14. Copyright Statement

Copyright (C) The Internet Society (2005). 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.

15. 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.

Pohl, Mark, Boschi

Expires August 2005

[Page 9]