IPFIX Working Group
Internet-Draft
Intended status: Informational
Expires: April 19, 2010

A. Kobayashi
NTT PF Lab.
B. Claise
Cisco Systems, Inc.
K. Ishibashi
NTT PF Lab.

October 16, 2009

IPFIX Mediation: Framework
draft-ietf-ipfix-mediators-framework-04

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79. This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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 April 19, 2010.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the

document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (http://trustee.ietf.org/license-info). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This document describes a framework for IPFIX Mediation. This framework details the IPFIX Mediation reference model and IPFIX Mediator components.

Table of Contents

1. Introduction

The IPFIX architectural components in [RFC5470] consist of IPFIX Devices and IPFIX Collectors communicating using the IPFIX protocol. In case of sustained growth of IP traffic in heterogeneous network environments, this Exporter-Collector architecture leads to the following problems: a lack of measurement system scalability, incompatibility of application requirements in a wide variety of measurement application, etc. These problems are described in detail in [IPFIX-MD-PS].

To acihive application requirements with limited system resources, IPFIX architecture needs to introduce an intermediate device between Exporters and Collectors. From a data manipulation point of view, this intermediate device provides the aggregation, correlation, filtering, and modification of Flow Records and/or PSAMP Packet Reports to optimize measurement system resources as a pre-process for the Collector. From a protocol conversion point of view, this intermediate device provides conversion into IPFIX or conversion of IPFIX transport protocols (e.g., from UDP to SCTP) to improve the reliability of the transport protocol.

This document introduces a generalized concept for such intermediate devices for IPFIX and describes the high-level architecture of IPFIX Mediation, key IPFIX Mediation architectural components, and characteristics of IPFIX Mediation.

This document is structured as follows: section 2 describes the terminology used in this document, section 3 gives an IPFIX/PSAMP

document overview, section 4 describes a high-level reference model, section 5 describes functional features related to IPFIX Mediation, section 6 describes combinations of components along with some application examples, section 7 describes consideration points of the encoding for IPFIX Message Headers, and section 8 describes the Information Elements used in an IPFIX Mediator.

2. Terminology and Definitions

The IPFIX-specific and PSAMP-specific terminology used in this document is defined in [RFC5101] and [RFC5476], respectively. The IPFIX Mediation-specific terminology used in this document is defined in [IPFIX-MD-PS]. However, as reading the problem statements document is not a prerequisite to reading this framework document, the definitions have been reproduced here along with additional definitions. In this document, as in [RFC5101] and [RFC5476], the first letter of each IPFIX-specific and PSAMP-specific term is capitalized along with the IPFIX Mediation-specific term defined here. The use of the terms "must", "should", and "may" in this document are informational only.

In this document, we use the generic term "record stream" to denote a set of flow- or packet-based data records with their additional information that flows from data sources, whether encoded in IPFIX protocol as IPFIX Data Records, or non-IPFIX protocols. In IPFIX protocol, we use the generic term Data Records for IPFIX Flow Records, PSAMP Packet Reports, and Data Records defined by Options Templates, unless an explicit distinction is required.

Transport Session Information

The Transport Session is specified in [RFC5101]. In SCTP, the Transport Session Information is the SCTP association. In TCP and UDP, the Transport Session Information corresponds to a 5-tuple {Exporter IP address, Collector IP address, Exporter transport port, Collector transport port, transport protocol}.

Original Exporter

An Original Exporter is an IPFIX Device that hosts the Observation Points where the metered IP packets are observed.

IPFIX Mediation

IPFIX Mediation is the manipulation and conversion of a record stream for subsequent export using the IPFIX protocol.

The following terms are used in this document to describe the architectural entities used by IPFIX Mediation.

Intermediate Process

An Intermediate Process takes a record stream as its input from Collecting Processes, Metering Processes, IPFIX File Readers, other Intermediate Processes, or other record sources; performs

some transformations on this stream, based upon the content of each record, states maintained across multiple records, or other data sources; and passes the transformed record stream as its output to Exporting Processes, IPFIX File Writers, or other Intermediate Processes, in order to perform IPFIX Mediation. Typically, an Intermediate Process is hosted by an IPFIX Mediator. Alternatively, an Intermediate Process may be hosted by an Original Exporter.

Specific Intermediate Processes are described below. However, this is not an exhaustive list.

Intermediate Conversion Process

An Intermediate Conversion Process is an Intermediate Process that

transforms non IPFIX into IPFIX, or manages the relation among Templates and states of incoming/outgoing transport sessions in the case of transport protocol conversion (e.g., from UDP to SCTP).

Intermediate Aggregation Process

An Intermediate Aggregation Process is an Intermediate Process that aggregates records based upon a set of Flow Keys or functions applied to fields from the record (e.g., binning and subnet aggregation).

Intermediate Correlation Process

An Intermediate Correlation Process is an Intermediate Process that adds information to records, noting correlations among them, or generates new records with correlated data from multiple records (e.g., the production of bidirectional flow records from unidirectional flow records).

Intermediate Selection Process

An Intermediate Selection Process is an Intermediate Process that selects records from a sequence based upon criteria-evaluated record values and passes only those records that match the criteria (e.g., filtering only records from a given network to a given Collector).

Intermediate Anonymization Process

An Intermediate Anonymization Process is an Intermediate Process that transforms records in order to anonymize them, to protect the identity of the entities described by the records (e.g., by

applying prefix-preserving pseudonymization of IP addresses).

IPFIX Mediator

An IPFIX Mediator is an IPFIX Device that provides IPFIX Mediation by receiving a record stream from some data sources, hosting one or more Intermediate Processes to transform that stream, and exporting the transformed record stream into IPFIX Messages via an

Exporting Process. In the common case, an IPFIX Mediator receives a record stream from a Collecting Process, but it could also receive a record stream from data sources not encoded using IPFIX, e.g., in the case of conversion from the NetFlow V9 protocol [RFC3954] to IPFIX protocol.

Specific types of IPFIX Mediators are defined below.

IPFIX Proxy

An IPFIX Proxy is an IPFIX Mediator that converts a record stream for the purpose of protocol conversion.

IPFIX Concentrator

An IPFIX Concentrator is an IPFIX Mediator that receives a record stream from one or more Exporters and performs correlation, aggregation, and/or modification.

IPFIX Distributor

An IPFIX Distributor is an IPFIX Mediator that receives a record stream from one or more Exporters and exports each record to one or more Collectors, deciding to which Collector(s) to export each record depending on the decision of an Intermediate Process.

IPFIX Masquerading Proxy

An IPFIX Masquerading Proxy is an IPFIX Mediator that receives a record stream from one or more Exporters to screen out parts of records according to configured policies in order to protect the privacy of the network's end users or to retain sensitive data of the exporting organization.

The following is a summary table for specific IPFIX Mediator types.

The abbreviation "IP" stands for Intermediate Process.

Table A: IPFIX Mediator Type Summary Table.

IPFIX Mediator Type	Number of hosted IPs	Intermediate Process Type
IPFIX Proxy	one or more	Intermediate Conversion Process
IPFIX Distributor	one or more	Intermediate Selection Process
IPFIX Concentrator	one or more	Intermediate Aggregation Process Intermediate Correlation Process
IPFIX Masquerading Proxy	one or more	Intermediate Anonymization Process

3. IPFIX/PSAMP Documents Overview

3.1. IPFIX Documents Overview

The IPFIX protocol [RFC5101] provides network administrators with access to IP flow information. The architecture for the export of measured IP flow information from an IPFIX Exporting Process to a Collecting Process is defined in [RFC5470], per the requirements defined in [RFC3917]. The IPFIX protocol [RFC5101] specifies how IPFIX Data Records and Templates are carried via a number of transport protocols from IPFIX Exporting Processes to IPFIX Collecting Processes. IPFIX has a formal description of IPFIX Information Elements, their names, types, and additional semantic information, as specified in [RFC5102]. [IPFIX-MIB] specifies the IPFIX Management Information Base. Finally, [RFC5472] describes what types 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. The storage of IPFIX Messages in a file is specified in [IPFIX-FILE].

3.2. PSAMP Documents Overview

The framework for packet selection and reporting [RFC5474] enables 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. The set of packet selection techniques (sampling and filtering) standardized by PSAMP is described in [RFC5475]. The PSAMP protocol [RFC5476] specifies the export of packet information from a PSAMP Exporting Process to a Collector. Like IPFIX, PSAMP has a formal description of its Information Elements, their names, types, and additional semantic information. The PSAMP information model is defined in [RFC5477]. [PSAMP-MIB] describes the PSAMP Management Information Base.

4. IPFIX Mediation Reference Model

The figure below shows the high-level IPFIX Mediation reference model based on [RFC5470]. This figure covers the various possible scenarios that can exist in an IPFIX measurement system.

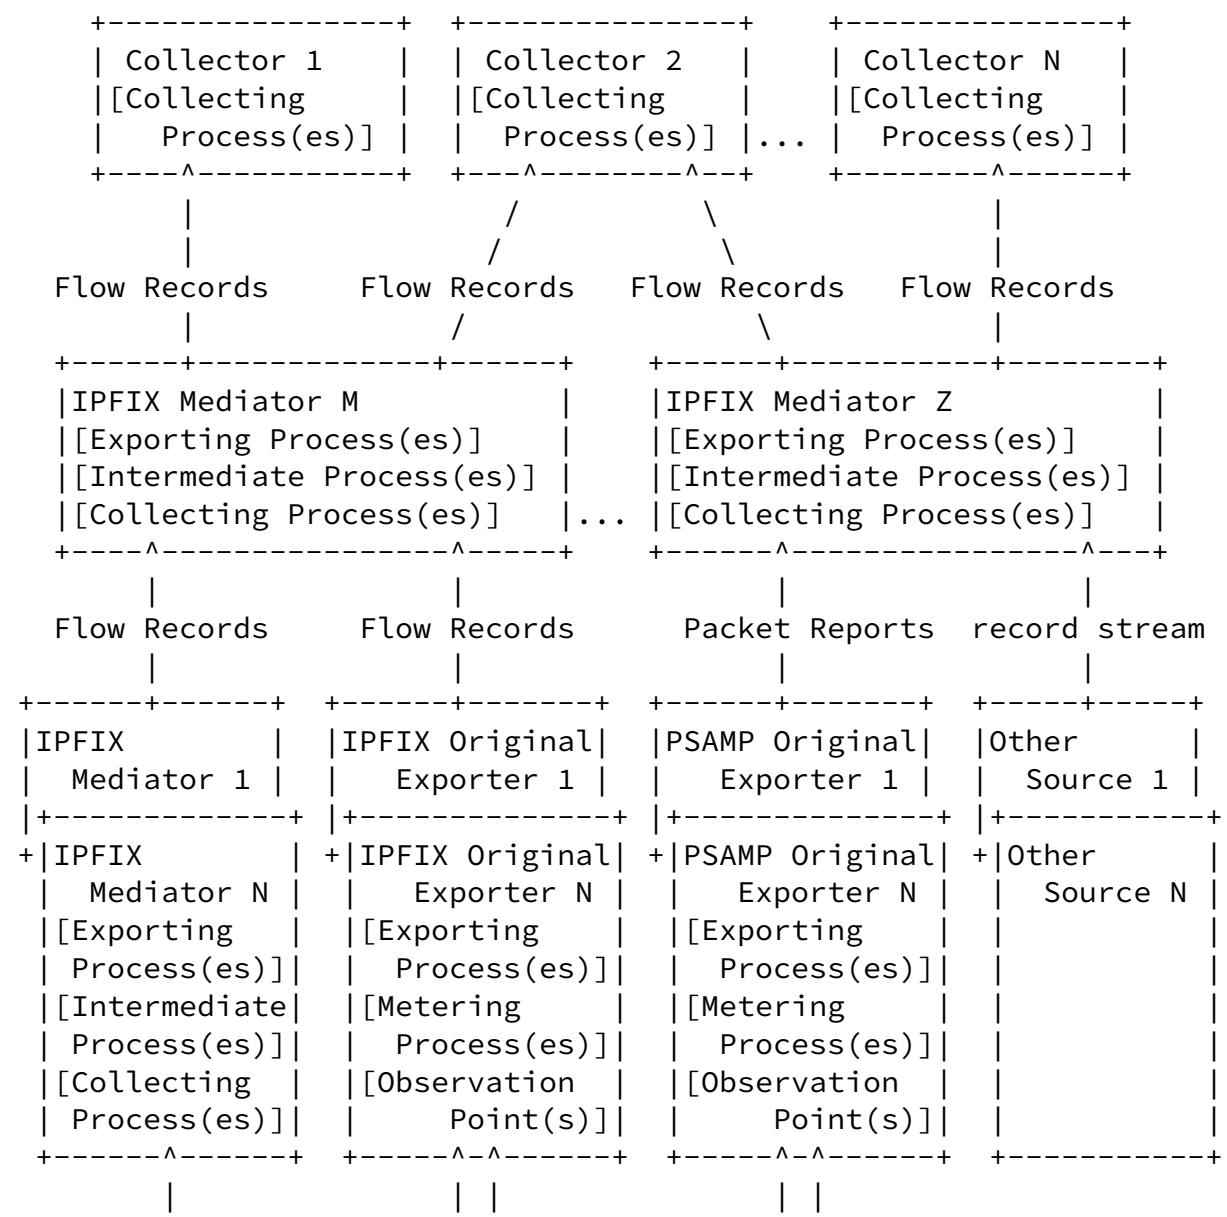

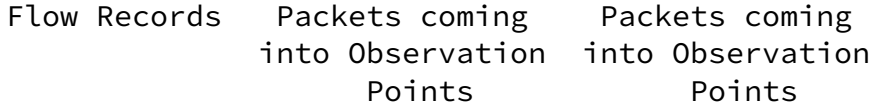

Figure A: IPFIX Mediation Reference Model Overview.

The functional components within each device are indicated within brackets []. An IPFIX Mediator receives IPFIX Flow Records or PSAMP Packet Records from other IPFIX Mediators, IPFIX Flow Records from

IPFIX Original Exporters, PSAMP Packet Reports from PSAMP Original Exporters, and/or a record stream from other sources. The IPFIX Mediator then exports IPFIX Flow Records and/or PSAMP Packet Reports to multiple Collectors and/or other IPFIX Mediators.

The figure below shows the basic IPFIX Mediator component model. Basically, an IPFIX Mediator, i.e., an IPFIX Proxy, IPFIX Masquerading Proxy, IPFIX Distributor, or IPFIX Concentrator, is composed of these components. An IPFIX Mediator contains one or more Intermediate Processes and one or more Exporting Processes. In typical case, it contains a Collecting Process, as described in the next figure.

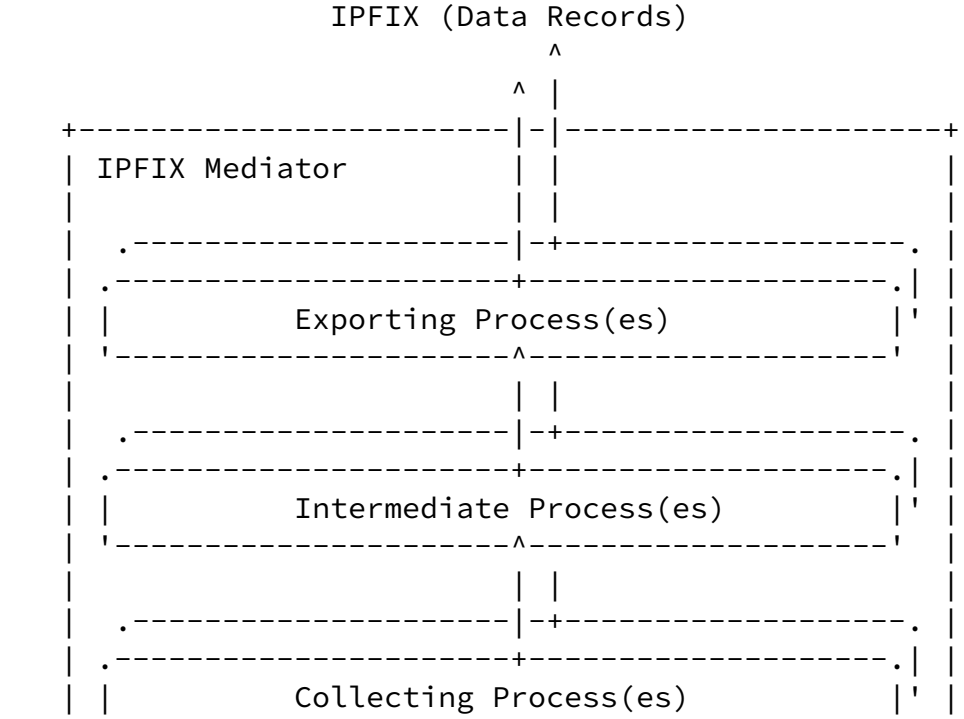

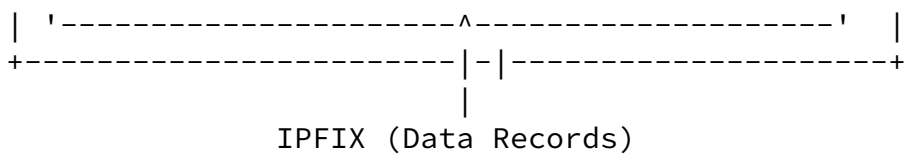

Figure B: Basic IPFIX Mediator Component Model.

However, other data sources are also possible: an IPFIX Mediator, i.e., IPFIX Proxy, can receive a record stream from non-IPFIX protocols such as NetFlow [RFC3954] exporter(s). This document does not make any particular assumption on how a record stream is transferred to an IPFIX Mediator. The figure below shows the IPFIX Mediator component model in the case of IPFIX protocol conversion from non-IPFIX exporters, so receiving a record stream.

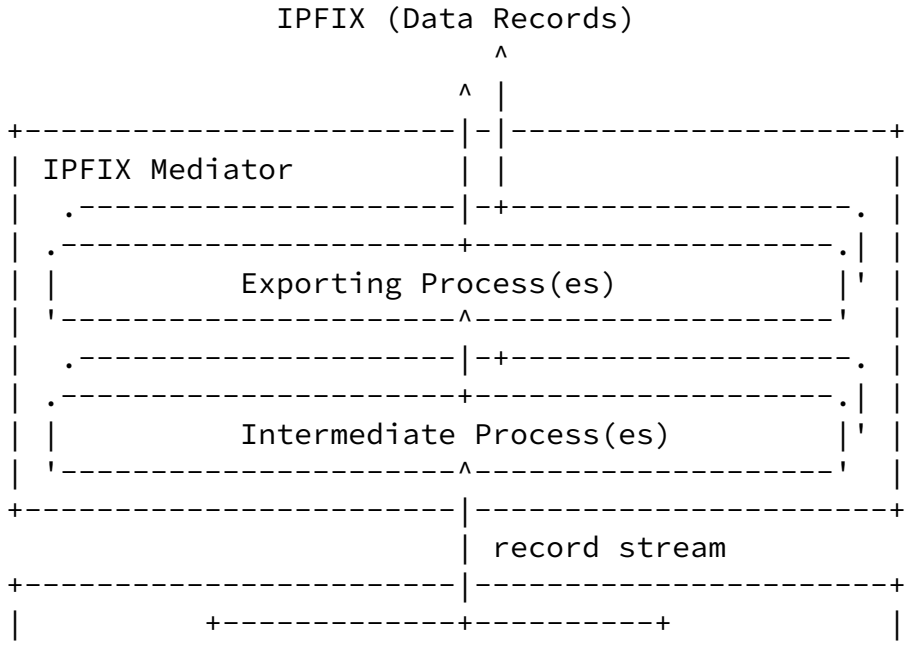

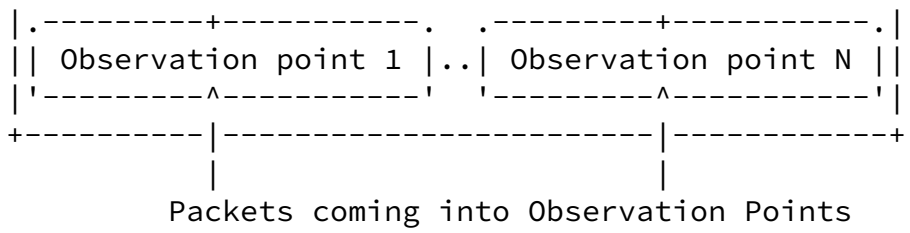

Figure C: IPFIX Mediator Component Model in IPFIX Protocol Conversion.

Alternatively, an Original Exporter may provide IPFIX Mediation by hosting one or more Intermediate Processes. The component model (Figure D) is composed by adding Intermediate Process(es) to the IPFIX Device model illustrated in [RFC5470]. In comparison with Figure 1 or 2 in [RFC5470], the Intermediate Process is located between IPFIX Metering Process(es), or PSAMP Metering Process(es) and Exporting Process(es).

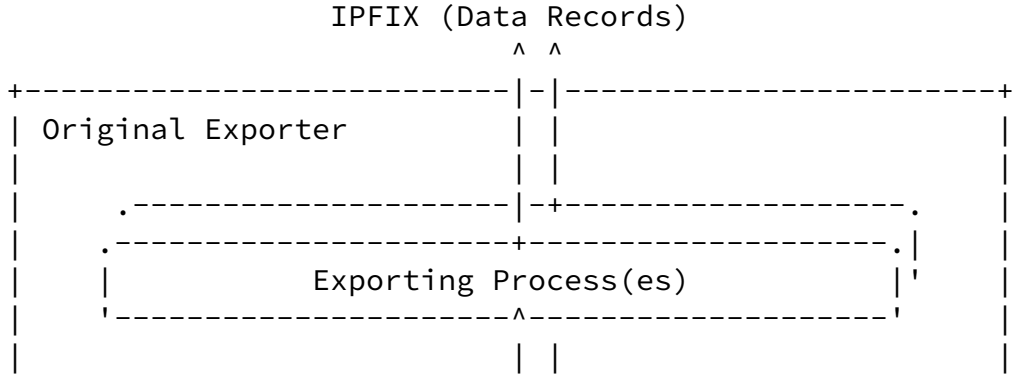

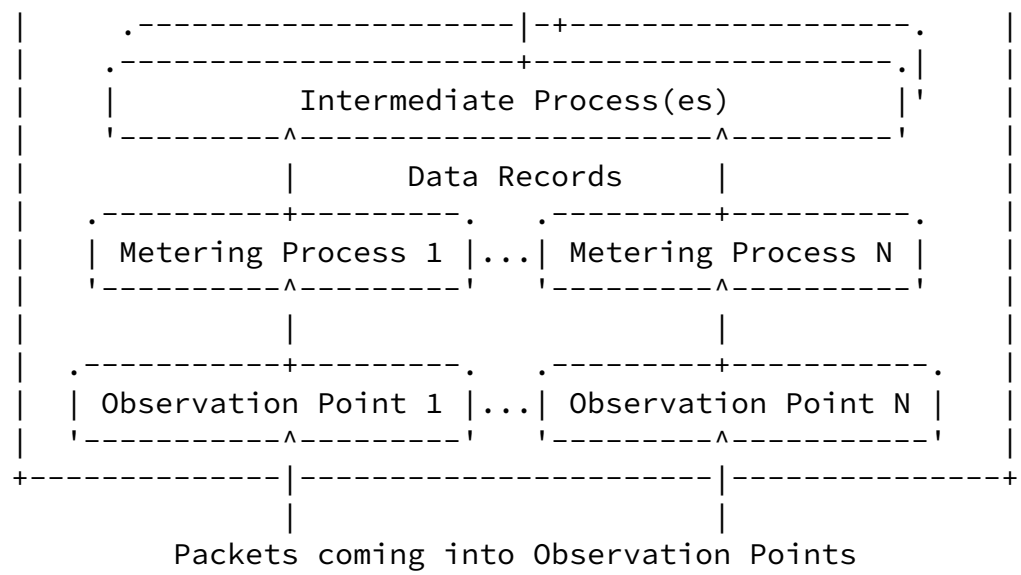

Figure D: IPFIX Mediation Component Model at Original Exporter.

In addition, an Intermediate Process may be collocated with an IPFIX File Reader and/or Writer. The following figure shows an IPFIX Mediation component model with an IPFIX File Writer and/or Reader.

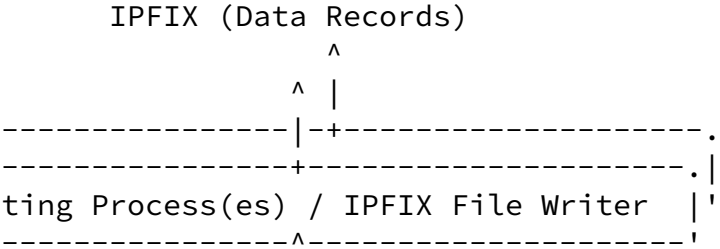

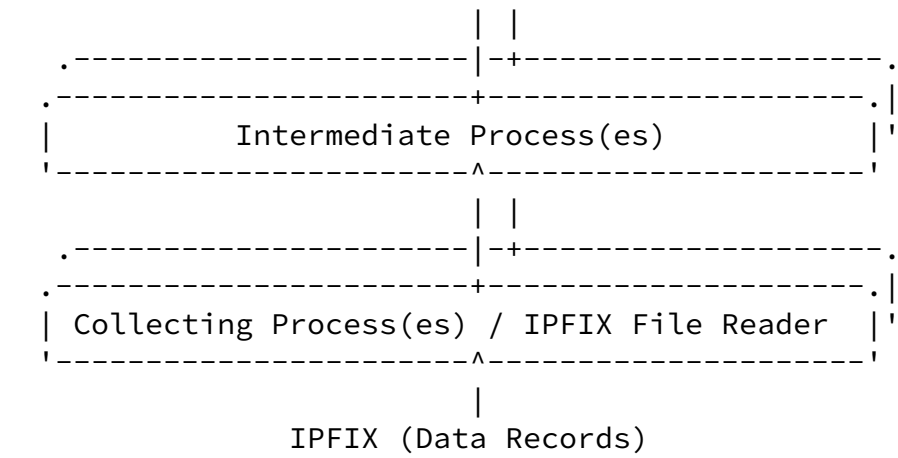

Figure E: IPFIX Mediation Component Model collocated with IPFIX File Writer/Reader.

5. IPFIX Mediation Functional Blocks

The following figure shows a functional block diagram of IPFIX Mediation in an IPFIX Mediator, having different Intermediate Process

types.

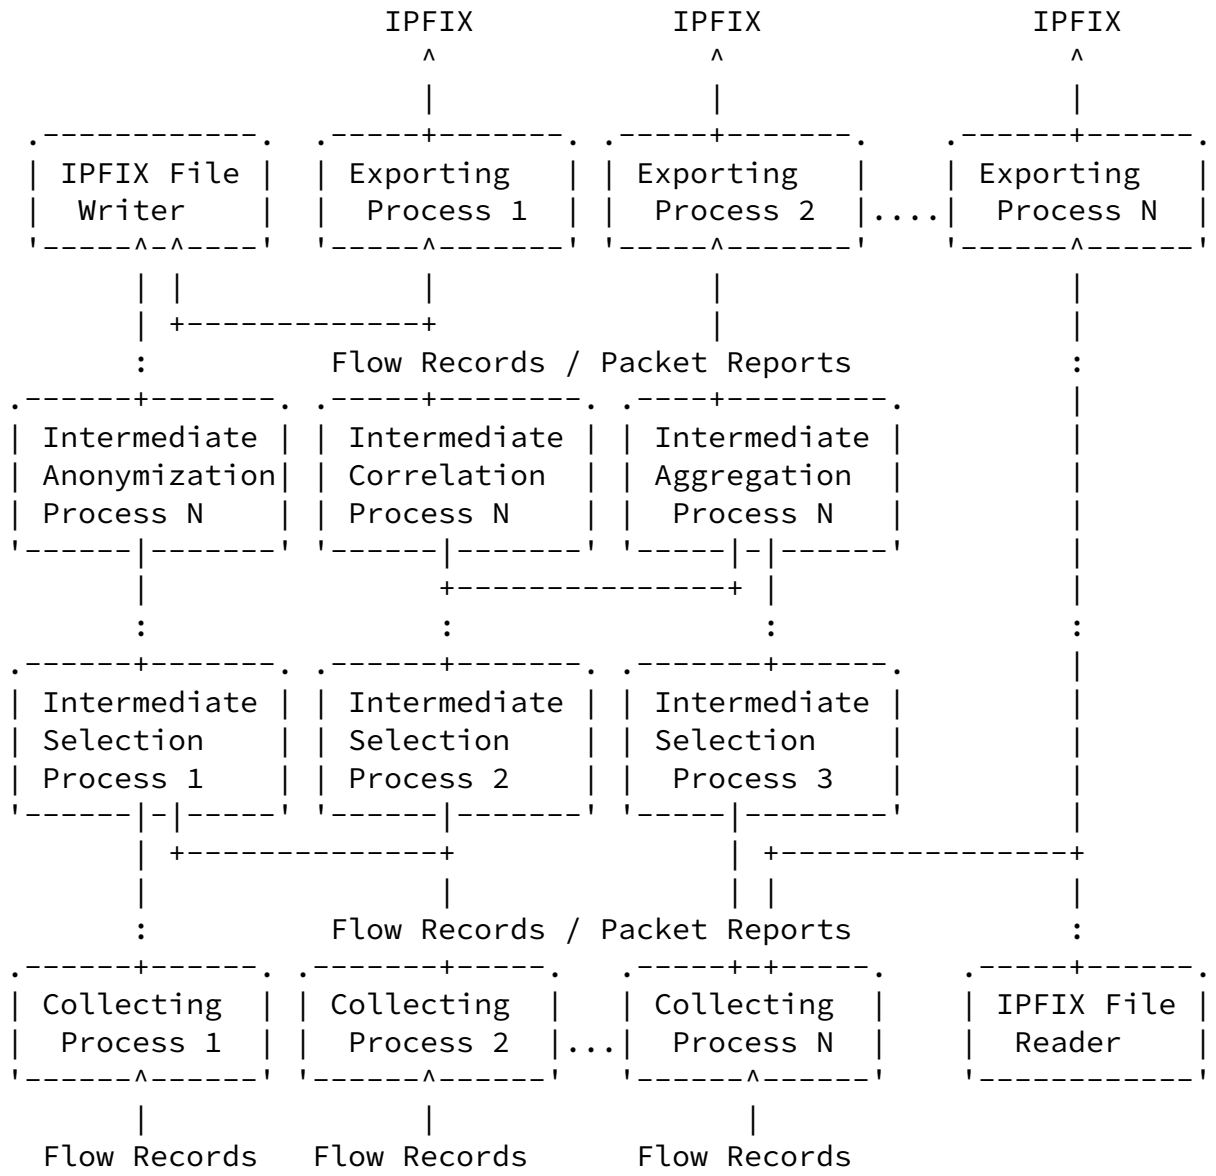

Figure F: IPFIX Mediation Functional Block.

5.1. Collecting Process

A Collecting Process in an IPFIX Mediator is not different than the Collecting Process described in [RFC5101]. Additional functions in an IPFIX Mediator include transmitting the set of Data Records and Control Information to one or more components, i.e., Intermediate Processes and other applications. In other words, a Collecting Process may duplicate the set and transmit it to one or more

components in sequence or in parallel. In the case of an IPFIX Mediator, the Control Information described in [RFC5470] includes IPFIX Message header information and Transport Session Information along with information about the Metering Process and the Exporting Process in an Original Exporter, e.g., sampling parameters.

5.2. Exporting Process

An Exporting Process in an IPFIX Mediator is not different than the Exporting Process described in [RFC5101]. Additional functions in an IPFIX Mediator may include the following.

- Receiving the trigger to transmit the Template Withdrawal Messages from Intermediate Process(es) when relevant Templates become invalid due to, for example, incoming session failure.

- Transmitting the measured data origins (e.g., Observation Point, Observation Domain ID, Original Exporter IP address, etc.) by encoding them to IPFIX. This function must be configurable.

5.3. Intermediate Process

An Intermediate Process is a key functional block for IPFIX Mediation. Its typical functions include the following:

- Generating a new record stream from an input record stream including context information (e.g., "Export Time", "Observation Domain ID", and Transport Session Information), and transmitting it to other components.

- Reporting statistics and interpretations for IPFIX Metering Processes, PSAMP Metering Processes, and Exporting Processes from an Original Exporter. See section 4 of [RFC5101] and section 6 of [RFC5476] for relevant statistics data structures and interpretations, respectively. This function must be configurable.

- Maintaining the configurable relation between Collecting Process(es)/Metering Process(es) and Exporting Process(es)/other Intermediate Process(es).

  A Collecting Process or Metering Process participating in IPFIX Mediation is associated with at least one Intermediate Process. Furthermore, an Intermediate Process is associated with at least one Exporting Process or another Intermediate Process. This relation can be configurable.

- Maintaining database(s) of all the Data Records in the case of an Intermediate Aggregation Process and an Intermediate Correlation Process. The function has the Data Record expiration rules described in the next subsection.

- Maintaining statistics on the Intermediate Process itself, such as the number of input/output Data Records, etc.

- Maintaining additional information about output record streams, which includes information related to Original Exporter, and some configuration parameters related to each function.

  In the case of an Intermediate Aggregation Process, Intermediate Anonymization Process, and Intermediate Correlation Process, the value of the "flowKeyIndicator" needs to be modified when modifying the data structure defined by an original Template.

5.3.1. Data Record Expiration

An Intermediate Aggregation Process and Intermediate Correlation Process need to have expiration conditions to export cached Data Records. In the case of the Metering Process in an Original Exporter, these conditions are described in [RFC5470]. In the case of the Intermediate Process, these conditions are as follows:

- If there are no input Data Records belonging to a cached Flow for a certain time period, aggregated Flow Records will expire. This time period should be configurable at the Intermediate Process.

- If the Intermediate Process experiences resource constraints, aggregated Flow Records may prematurely expire (e.g., lack of memory to store Flow Records).

- For long-running Flows, the Intermediate Process should cause the Flow to expire on a regular basis or on the basis of an expiration policy. This periodicity or expiration policy should be configurable at the Intermediate Process.

In the case of an Intermediate Correlation Process, a cached Data Record may be prematurely expired (and discarded) when no correlation

can be computed with newly received Data Records. For example, an Intermediate Correlation Process computing one way delay may discard the cached Packet Report when no other matching packet Report are observed within a certain time period.

5.3.2. Specific Intermediate Processes

This section shows the functional blocks of specific Intermediate Processes.

5.3.2.1. Intermediate Conversion Process

When receiving a non IPFIX record stream, the Intermediate Conversion Process covers the following functions:

- Retrieving the value for each Information Element from each record, and converting the Information Element IDs (e.g., from NetFlow V9 protocol [RFC3954] to IPFIX Information Model [RFC5102]).
- Transforming a record stream into Data Records, (Options) Template Records, and/or Data Records defined by Options Templates.
- Transforming additional information (e.g., Sampling rate, Sampling algorithm, and observation information) into Data Records or Data Records defined by Options Templates.

In the case of IPFIX transport protocol conversion, the Intermediate Conversion Process covers the following functions.

- Relaying Data Records, (Options) Template Records, and Data Records defined by Options Templates.
- Setting the trigger for the Exporting Process in order to export IPFIX Template Withdrawal Messages relevant to the Templates when Templates becomes invalid due to, for example, incoming session failure.

- Maintaining the mapping information about Transport Sessions, Observation Domain IDs, and Template IDs on the incoming/outgoing sides to confirm the appropriateness of the scope field values in Data Records defined by Options Templates and of IPFIX Template Withdrawal Messages.

5.3.2.2. Intermediate Selection Process

An Intermediate Selection Process has analogous functions to the PSAMP Selection Process described in [RFC5475]. The difference is that the Intermediate Selection Process takes a record stream, e.g., Flow Records or Packet Reports, rather than observed packets.

The typical function is property match filtering that retrieves a record stream of interest. The function selects a Data Record if the

value of a specific field in the Data Record equals a configured value or falls within a configured range.

5.3.2.3. Intermediate Aggregation Process

An Intermediate Aggregation Process covers the following typical functions:

- Merging a set of Data Records within a certain time period into one Flow Record by adding up the counters.

- Maintaining statistic and additional information about aggregated Flow Records.

  The statistics for an aggregated Flow Record may include the number of original Data Records and the maximum and minimum values of per-flow counters. Additional information may include a certain time period, a new set of Flow Keys, and observation location information involved in the Flow aggregation. Observation location information can be tuples of (Observation Point, Observation Domain ID, Original Exporter IP address) or another identifier indicating an area.

- Data Records aggregation, which can be done in the following ways:

* Spatial composition

  With spatial composition, Data Records sharing common properties are merged into one Flow Record within a certain time period. One typical aggregation can be based on a new set of Flow Keys. Generally, a shorter set of common properties than an original set of Flow Keys creates more aggregated Flow Records. Another aggregation can be based on a set of Observation Points within an Observation Domain, on a set of Observation Domains within an Exporter, or on a set of Exporters.

  If some fields do not serve as Flow Keys, the Intermediate Aggregation Process determines these values by the first received Data Record, a specific Exporter IP address, or other arbitrary decision. The Intermediate Aggregation Process should be consistent in its decision method in an Intermediate Aggregation Process.

  Furthermore, a new identifier indicating a group of observation locations can be introduced, for example, to indicate an area on a large network or a link aggregation interface composed of

  physical interfaces, or a set of values of a specified field in the original Data Records.

* Temporal composition

  With temporal composition, Flow Records on one Observation Point are merged into one Flow Record within a certain time period. For example, short-period Flow Records, e.g., 1 minute, are merged into a long-period Flow Record, e.g., 30 minutes. In case of a long-running Flow, The temporal composition provides some advantages:

  + Reducing the number of Flow Records for long-running Flows, reducing the export bandwidth requirements.

  + Computing the real active time period for long-running Flows by summing up the short-period Flow Records.

  + Producing more precise maximum and minimum values without

increasing the number of Flow Records on a Collector.

For example, short-period Flow Records created at a Metering Process by configuring a short active time, e.g., 1 or 10 sec, are merged within a certain time period, e.g., 60 or 300 sec at an Intermediate Aggregation Process. While merging, new counters, such as the maximum and minimum, can be created.

When some traffic requires timely traffic monitoring and other traffic does not, a combination of the Intermediate Selection Process and Intermediate Aggregation Process is useful, as described in section 6.

5.3.2.4. Intermediate Anonymization Process

An Intermediate Anonymization Process covers the following typical functions.

- Deleting specified fields

  The function deletes existing fields in accordance with some instruction rules. Examples include hiding network topology information and private information. In the case of feeding Data Records to end customers, disclosing vulnerabilities is avoided by deleting fields, e.g., "ipNextHopIP{v4|v6}Address", "bgpNextHopIP{v4|v6}Address", "bgp{Next|Prev}AdjacentAsNumber", and "mplsLabelStackSection", described in [RFC5102].

- Anonymizing value of specified fields

  The function modifies the value of specified fields. Examples include anonymizing customers' private information, such as IP address and port number, in accordance with a privacy protection policy. The Intermediate Anonymization Process may also report anonymized fields and the anonymization method as additional information.

5.3.2.5. Intermediate Correlation Process

An Intermediate Correlation Process can be viewed as a special case

of the Intermediate Aggregation Process, covering the following typical functions:

- Producing new information including metrics, counters, attributes, or packet property parameters by evaluating the correlation among sets of Data Records or among Data Records and other meta data after gathering sets of Data Records within a certain time period.

- Adding new fields into a Data Record or creating a new Data Record.

A correlation for Data Records can be done in the following ways.

- One-to-one correlation between Data Records, with the following examples:

  * One-way delay, Packet delay variation in [RFC5481]
    The metrics follow from the correlation of the timestamp value on a pair of Packet Reports indicating an identical packet from different Exporters.

  * Packet inter-arrival time or jitter
    The metrics follow from the correlation of the timestamp value on consecutive Packet Reports from a single Exporter.

  * Rate-limiting ratio, compression ratio, optimization ratio, etc.
    The data values follow from the correlation of Data Records indicating identical a Flow observed on the incoming/outgoing points of a WAN interface.

- Correlation amongst Data Records, with the following examples:

  * Bidirectional Flow composition
    The method of exporting and representing a bidirectional flow (Biflow) is described in [RFC5103]. The bidirectional flow

composition is a special case of Flow Key aggregation. The Flow Records are merged into one Flow Record as Biflow, if Non-directional Key Fields matches and the Directional Key Field matches their reverse direction counterparts. The direction assignment method to assign the Biflow Source and Destination

as additional information may be reported. In the case of an Intermediate Aggregation Process, the direction may be assigned as "arbitrary".

* Average/maximum/minimum for packets, bytes, one-way delay, packet loss, etc.
The data values follow from the correlation of multiple Data Records while the Intermediate Aggregation Process executes.

o Correlation between Data Record and other meta data

Typical examples are derived packet property parameters described in [RFC5102]. The parameters are retrieved based on the value of the specified field in an input Data Record, compensating for traditional exporting devices or probes that are unable to add packet property parameters. Therefore, Collectors do not need to recognize the differences among implementations of routers from several vendors or among Exporter types, such as router, switch, or probe. Typical derived packet property parameters are as follows:

* "bgpNextHop{IPv4|IPv6}Address" described in [RFC5102], which indicates the egress router of a network domain. It is useful for making a traffic matrix that covers the whole network domain.

* BGP Communities attribute
This attribute indicates tagging for routes of geographical and topological information and source types (e.g., transit, peer, or customer) as described in [RFC4384]. Therefore, network administrators can monitor the geographically-based or source type-based traffic volume by correlating the attribute.

* "mplsVpnRouteDistinguisher" described in [RFC5102]
This value indicates the VPN customer's identification, which cannot be extracted from the core router in MPLS networks. Therefore, network administrators can monitor the customer-based traffic volume on even core routers.

6. Component Combination

An IPFIX Mediator should be able to simultaneously support more than one Intermediate Process. Each of the Intermediate Processes should be independent. Multiple Intermediate Processes generally are configured in the following ways.

- Parallel Intermediate Processes

  To feed a record stream to different applications having different requirements, the Intermediate Processes are located in parallel.

- Serial Intermediate Processes

  To execute flexible manipulation of a record stream, the Intermediate Processes are connected serially. In that case, an output record stream from one Intermediate Process forms an input record stream for a succeeding Intermediate Process.

In consideration of resource contention, the series of Intermediate Processes and associated Exporting Processes has preferably the same priority value determined based on application requirements.

In addition to the combination of Intermediate Processes, the combination of some components (Exporting Process, Collecting Process, IPFIX File Writer and Reader) can be applied to provide various data reduction techniques. This section shows some combinations along with examples.

6.1. Data-based Collector Selection

The combination of one or more Intermediate Selection Processes and Exporting Processes can determine to which Collector input Data Records are exported. Applicable examples include exporting Data Records to a dedicated Collector on the basis of customer or organization. For example, an Intermediate Selection Process selects Data Records from record stream duplicated in a Collecting Process on the basis of the peering autonomous system number, and an Exporting Process sends them to a dedicated Collector, as shown in the following figure.

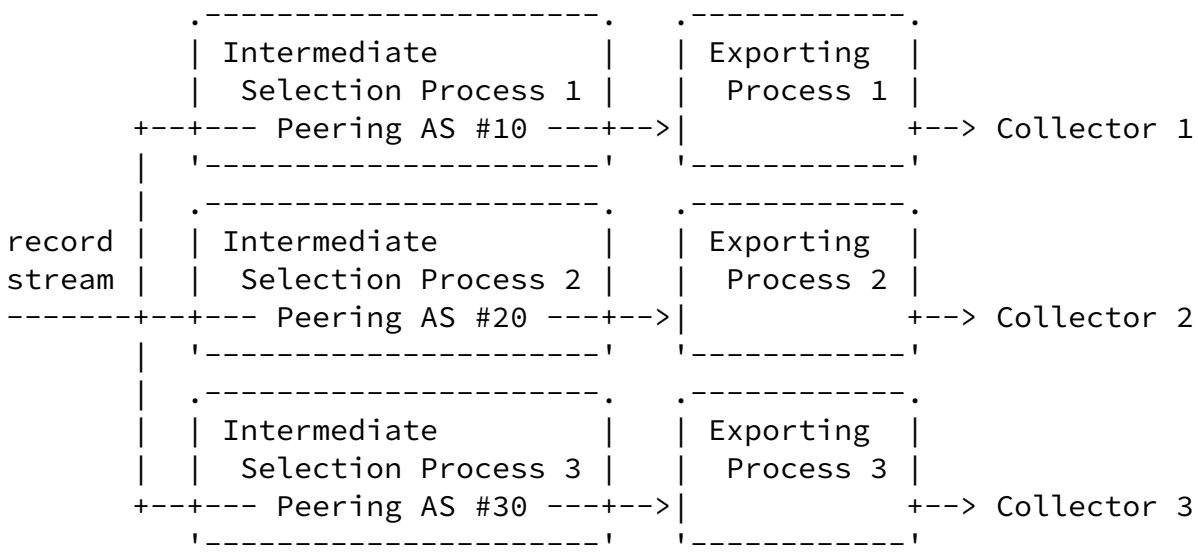

Figure G: Data-based Collector Selection.

6.2. Flow Selection and Aggregation

The combination of one or more Intermediate Selection Processes and Intermediate Aggregation Processes can efficiently reduce the amount of Flow Records. The combination structure is similar to the concept of the Composite Selector described in [RFC5474]. For example, an Intermediate Selection Process selects Flows consisting of a small number of packets and then transmits them to an Intermediate Aggregation Process. Another Intermediate Selection Process selects other Flow Records and then transmits them to an Exporting Process, as shown in the following figure. This results in aggregation on the basis of the distribution of the number of packets per Flow.

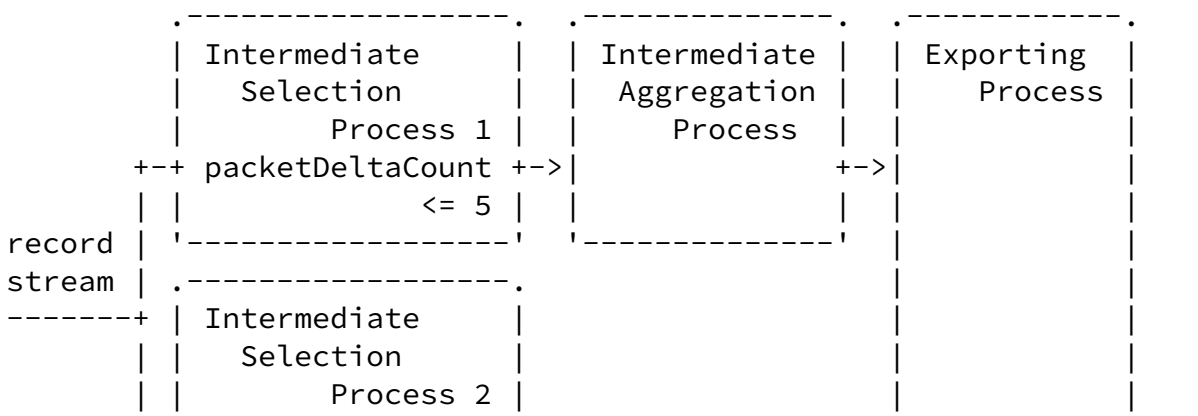

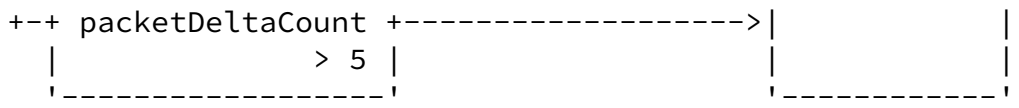

Figure H: Flow Selection and Aggregation Example.

6.3. IPFIX File Writer/Reader

The IPFIX File Writer/Reader on an IPFIX Mediator complies with [IPFIX-FILE]. The IPFIX File Writer stores Data Records in a file system. When Data Records include problematic Information Elements, an Intermediate Anonymization Process can delete these fields before the IPFIX File Writer handles them, as shown in the following figure.

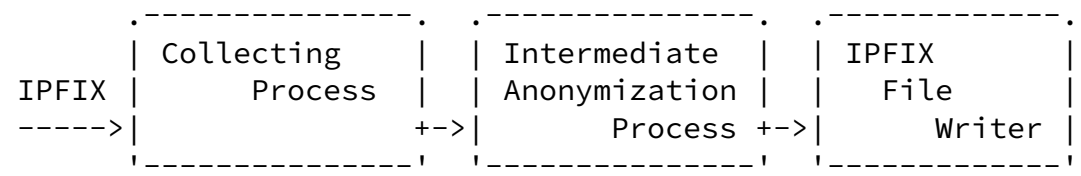

Figure I: IPFIX Mediation Example with IPFIX File Writer.

In contrast, the IPFIX File Reader retrieves stored Data Records when administrators want to retrieve past Data Records from a given time period. If the data structure of the Data Records from the IPFIX File Reader is different from what administrators want, an Intermediate Anonymization Process and Intermediate Correlation Process can modify the data structure, as shown in the following figure.

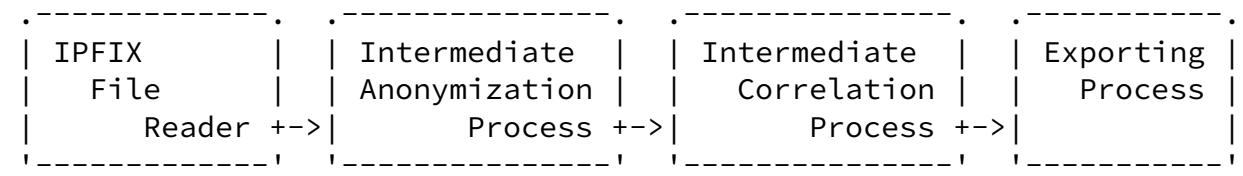

Figure J: IPFIX Mediation Example with IPFIX File Reader.

7. Encoding for IPFIX Message Header

The IPFIX Message Header [RFC5101] includes Export Time, Sequence Number, and Observation Domain ID fields. This section describes some consideration points for the IPFIX Message Header encoding.

Export Time

An IPFIX Mediator can set the Export Time in two ways.

* Case 1: keeping the field value of incoming Transport Sessions

* Case 2: setting the time at which an IPFIX Message leaves the IPFIX Mediator

In case 2, the IPFIX Mediator needs to handle any delta time stamp fields, such as "flowStartDeltaMicroseconds" and "flowEndDeltaMicroseconds", described in [RFC5102].

Sequence Number

In the case of an IPFIX Proxy relaying a one-to-one Transport Session, the IPFIX Proxy needs to handle the Sequence Number value when the incoming Transport Session shuts down and starts.

Observation Domain ID

An IPFIX Mediator can set the Observation Domain ID independently

of the incoming Observation Domain ID. There are two consideration points.

* Case 1: relaying an IPFIX Message after replacing each incoming Observation Domain ID with a new value in the case of an IPFIX Proxy and an IPFIX Concentrator

* Case 2: aggregating incoming Flow Records in the case of an IPFIX Concentrator

In case 1, an IPFIX Proxy needs to set the appropriate scope fields in Data Records defined in Options Template Records when the incoming Observation Domain IDs are used as the scope fields. In case 2, according to the description in [RFC5101], an IPFIX Concentrator needs to set a value of 0 for the Observation Domain ID. In that case, the IPFIX Concentrator can add a new field to the Flow Record instead of the Observation Domain ID. The field indicates the largest set of Observation Points for an aggregated Flow Record.

8. Information Model

IPFIX Mediation reuses the general information models from [RFC5102] and [RFC5477]. However, several Intermediate Processes would potentially require additional Information Elements as follows:

- Number of original Data Records belonging to output aggregated Flow Records as described in section 5.3.2.3. Something similar to the "Flow" Information Element, id 3 in NetFlow version 9 [RFC3954].

- New observation domain information instead of Observation Domain ID in IPFIX Concentrator as described in section 7.

- Maximum and minimum values for packet count and octet count as described in the "time composition" paragraph in section 5.3.2.3.

- Some metrics related to network performance, e.g., one-way delay, packet inter-arrival time, etc., as described in section 5.3.2.5.

- Anonymization method and report on the anonymized fields as

described in section 5.3.2.4.

- Report on the applied treatment items in IPFIX Mediation.

9. Security Considerations

An IPFIX measurement system must also prevent the security threats related to IPFIX Mediation that follow as well as the security threats described in the security consideration section in [RFC5101].

- Attacks against IPFIX Mediators

  IPFIX Mediators need to prevent unauthorized access or denial-of-service (DoS) attacks. One solution is for IPFIX Mediators to host the packet filter function to reject malicious packets at an outside interface.

- Man-in-the-middle attacks by untrusted IPFIX Mediators

The Exporter-Mediator-Collector structure model would increase the risk of man-in-the-middle attacks. One solution is that IPFIX Collectors and Exporters must verify trusted IPFIX Mediators to prevent connection to untrusted IPFIX Mediators.

- Configuration of IPFIX Mediation

  In the case of IPFIX Distributors and IPFIX Masquerading Proxies, an accidental misconfiguration and unauthorized access to configuration data could lead to the crucial problem of disclosure of confidential traffic data.
  To eliminate these risks, IPFIX Mediators must provide the authentication function for authorized administrators and the facilities to help in tracing configuration changes to their origins.

10. IANA Considerations

This document has no actions for IANA.

11. References

11.1. Normative References

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

[RFC5476] Claise, B., Quittek, J., and A. Johnson, "Packet Sampling (PSAMP) Protocol Specifications", March 2009.

11.2. Informative References

[IPFIX-FILE] Trammell, B., Boschi, E., Mark, L., Zseby, T., and A. Wagner, "An IPFIX-Based File Format", draft-ietf-ipfix-file-05 (work in progress) , August 2009.

[IPFIX-MD-PS] Kobayashi, A., Claise, B., Nishida, H., Sommer, C., Dressler, F., and E. Stephan, "IPFIX Mediation: Problem Statement", draft-ietf-ipfix-mediation-problem-statement-05 (work in progress) , July 2009.

[IPFIX-MIB] Dietz, T., Claise, B., and A. Kobayashi, "Definitions of Managed Objects for IP Flow Information Export", draft-ietf-ipfix-mib-07 (work in progress) , July 2009.

[PSAMP-MIB] Dietz, T. and B. Claise, "Definitions of Managed Objects for Packet Sampling", draft-ietf-psamp-mib-06 (work in progress) , June 2006.

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

[RFC3954] Claise, B., "Cisco Systems NetFlow Services Export Version 9", October 2004.

[RFC4384] Meyer, D., "BGP Communities for Data Collection", February 2006.

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

[RFC5103] Trammell, B. and E. Boschi, "Bidirectional Flow Export Using IP Flow Information Export (IPFIX)", January 2008.

[RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", March 2009.

[RFC5472] Zseby, T., Boschi, E., Brownlee, N., and B. Claise, "IPFIX Applicability", March 2009.

[RFC5474] Duffield, N., "A Framework for Packet Selection and Reporting", March 2009.

[RFC5475] Zseby, T., Molina, M., Duffield, N., Niccolini, S., and F. Raspall, "Sampling and Filtering Techniques for IP Packet Selection", March 2009.

[RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", March 2009.

[RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", March 2009.

Appendix A. Acknowledgements

We would like to thank the following persons: Gerhard Muenz for the thorough detail review and significant contribution regarding the improvement of whole sections; Daisuke Matsubara, Tsuyoshi Kondoh, Hiroshi Kurakami, Haruhiko Nishida for contribution during the initial phases of the document; Brian Trammel for contribution regarding the improvement of terminologies section; Nevil Brownlee, Juergen Quittek for the technical reviews and feedback.

Authors' Addresses

Atsushi Kobayashi
NTT Information Sharing Platform Laboratories
3-9-11 Midori-cho
Musashino-shi, Tokyo 180-8585
Japan

Phone: +81-422-59-3978
Email: akoba@nttv6.net

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
Diegem 1831
Belgium

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

Keisuke Ishibashi
NTT Information Sharing Platform Laboratories
3-9-11 Midori-cho
Musashino-shi 180-8585
Japan

Phone: +81-422-59-3978

Email: ishibashi.keisuke@lab.ntt.co.jp