

IPFIX Working Group
Internet-Draft
Intended Status: Standards Track
Expires: April 25, 2011

B. Claise
Cisco Systems, Inc.
A. Kobayashi
NTT PF Lab.
B. Trammell
Hitachi Europe
October 25, 2010

Specification of the Protocol for IPFIX Mediations
draft-claise-ipfix-mediation-protocol-02

Abstract

This document specifies the IP Flow Information Export (IPFIX) protocol specific to the Mediation.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and 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/lid-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, 2011.

Copyright Notice

Copyright (c) 2010 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents

<Claise, et. Al>

Expires April 25 2011

[Page 1]

Internet-Draft <Protocol for IPFIX Mediations> October 2010
carefully, as they describe your rights and restrictions with
respect to this document. Code Components extracted from this
document must include Simplified BSD License text as described
in Section 4.e of the Trust Legal Provisions and are provided
without warranty as described in the Simplified BSD License.

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL",
"SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY",
and "OPTIONAL" in this document are to be interpreted as
described in RFC 2119 [RFC2119].

Table of Contents

1. Introduction.....	3
1.1. IPFIX Documents Overview.....	4
1.2. IPFIX Mediator Documents Overview.....	4
1.3. Relationship with IPFIX and PSAMP.....	5
2. Terminology.....	6
3. Specifications.....	9
3.1. Encoding of IPFIX Message Header.....	9
3.2. Template Management.....	11
3.2.1. Template Management Without Template Record Change	11
3.2.2. Template Management With Template Record Change..	12
3.3. Time Management.....	12
3.4. Observation Point Management.....	13
3.4.1. Observation Domain Management.....	15
3.5. Specific Reporting Requirements.....	15
3.5.1. The Flow Keys Options Template.....	16
3.5.2. IPFIX Protocol Options Template.....	16
3.5.3. IPFIX Mediator Options Template.....	17
3.6. The Collecting Process's Side.....	17
3.7. Sampling Management.....	18
4. New Information Elements.....	18
5. Security Considerations.....	18
5.1. Avoiding Security Downgrade.....	19
5.2. End-to-End Assertions for Mediators.....	19
6. IANA Considerations.....	20
6.1. originalExporterIPv4Address.....	21
6.2. originalExporterIPv6Address.....	21
7. References.....	21
7.1. Normative References.....	21
7.2. Informative References.....	22

Internet-Draft	<Protocol for IPFIX Mediations>	October 2010
8. Author's Addresses.....		23
9. Appendix A. Additions to XML Specification of IPFIX		
Information Elements.....		24

TO DO

- See the EDITOR'S NOTE within the document
- Check that all documents: this one, [IPFIX-MED-ANON], [IPFIX-MED-FLOWSEL], [IPFIX-MED-AGGR] all fit together.

1. Introduction

The IPFIX architectural components in [RFC5470] consist of IPFIX Devices and IPFIX Collectors communicating using the IPFIX protocol [RFC5101], which specifies how to export IP Flow information. 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 relevant Information Element is specified in the IPFIX Information Model [RFC5102], registered with IANA, or specified as an enterprise-specific Information Element. The specifications in the IPFIX protocol [RFC5101] have not been defined in the context of an IPFIX Mediator receiving, aggregating, correlating, anonymizing, etc... Flow Records from the one or multiple Exporters. Indeed, the IPFIX protocol must be adapted for Intermediate Processes, as defined in the IPFIX Mediation Reference Model as specified in the Figure A of [IPFIX-MED-FMWK], which is based on the IPFIX Mediation Problem Statement [RFC5982].

This document specifies the IP Flow Information Export (IPFIX) protocol in the context of the implementation and deployment of IPFIX Mediators. The use of the IPFIX protocol within a Mediator -- a device which contains both as an Exporting Process and a Collecting Process -- has an impact on the technical details of the usage of the protocol. An overview of the technical problem is covered in section 6 of the [RFC5982]: loss of original exporter information, loss of base time information, transport sessions management, loss of Options Template Information,

Internet-Draft <Protocol for IPFIX Mediations> October 2010
Template Id management, considerations for network topology,
and IPFIX Mediation interpretation, and considerations for
aggregation.

The specifications in this document are based on the IPFIX
protocol specifications but adapted according to the IPFIX
Mediation Framework [IPFIX-MED-FMWK].

1.1. 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 out of an IPFIX Exporting Process to a Collecting
Process is defined in the IPFIX Architecture [RFC5470], per
the requirements defined in RFC 3917 [RFC3917].

The IPFIX Architecture [RFC5470] specifies how IPFIX Data
Records and Templates are carried via a congestion-aware
transport protocol from IPFIX Exporting Processes to IPFIX
Collecting Processes.

IPFIX has a formal description of IPFIX Information Elements,
their name, type and additional semantic information, as
specified in the IPFIX Information Model [RFC5102].

The IPFIX Applicability Statement [RFC5472] 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.

"IPFIX Mediation: Problem Statement" [RFC5982], describing the
IPFIX Mediation applicability examples, along with some problems
that network administrators have been facing, is the basis for
the "IPFIX Mediation: Framework" [IPFIX-MED-FMWK]. This
framework details the IPFIX Mediation reference model and the
components of an IPFIX Mediator.

1.2. IPFIX Mediator Documents Overview

The "IPFIX Mediation: Problem Statement" [RFC5982] provides an
overview of the applicability of Mediators, and defines
requirements for Mediators in general terms. This document is

Internet-Draft <Protocol for IPFIX Mediations> October 2010
of use largely to define the problems to be solved through the
deployment of IPFIX Mediators, and to provide scope to the role
of Mediators within an IPFIX collection infrastructure.

The "IPFIX Mediation: Framework" [IPFIX-MED-FMWK] provides more
architectural details of the arrangement of Intermediate
Processes within a Mediator.

The details of specific Intermediate Processes, when these have
additional export specifications (e.g., metadata about the
intermediate processing conveyed through IPFIX Options
Templates), are each treated in their own document (e.g., the
"IP Flow Anonymisation Support" [IPFIX-MED-ANON]). Documents
specifying the operations of specific Intermediate Processes
cover the operation of these Processes within the Mediator
framework, and complying to the specifications given in this
document; they may additionally specify the operation of the
process independently, outside the context of a Mediator, when
this is appropriate. As of today, these documents are:

1. "IP Flow Anonymisation Support", [IPFIX-MED-ANON], which
describes anonymisation techniques for IP flow data and the
export of anonymised data using the IPFIX protocol.
2. "Flow Selection Techniques" [IPFIX-MED-FLOWSEL], which
described the process of selecting a subset of flows from all
flows observed at an observation point, along with the
motivations, and some specific flow selection techniques.
3. "Exporting Aggregated Flow Data using the IP Flow Information
Export" [IPFIX-MED-AGGR] which describes Aggregated Flow export
within the framework of IPFIX Mediators and defines an
interoperable, implementation-independent method for Aggregated
Flow export.

1.3. Relationship with IPFIX and PSAMP

The specification in this document applies to the IPFIX
protocol specifications [RFC5101]. All specifications from
[RFC5101] apply unless specified otherwise in this document.

As the Packet Sampling (PSAMP) protocol specifications
[RFC5476] are based on the IPFIX protocol specifications, the
specifications in this document are also valid for the PSAMP
protocol. Therefore, the method specified by this document
also applies to PSAMP.

The IPFIX-specific and PSAMP-specific terminology used in this document is defined in [RFC5101] and [RFC5476], respectively. The IPFIX Mediation terms related to the aggregation, such as the Interval, Aggregated Flow, and Aggregated Function are defined in [IPFIX-MED-AGGR].

The IPFIX Mediation-specific terminology used in this document is defined in "IPFIX Mediation: Problem Statement" [RFC5982], and reuse in "IPFIX Mediation: Framework" [IPFIX-MED-FMWK]. However, since those two documents are an informational RFC, the definitions have been reproduced here along with additional definitions.

Similarly, since the [IPFIX-MED-ANON] is an experimental RFC, the Anonymisation Record, Anonymised Data Record, and Intermediate Anonymisation Process terms, specified in [IPFIX-MED-ANON], are also reproduced here.

In this document, as in [RFC5101], [RFC5476], [IPFIX-MED-AGGR], and [IPFIX-MED-ANON], the first letter of each IPFIX-specific and PSAMP-specific term is capitalized along with the IPFIX Mediation-specific term defined here. In this document, we call "record stream" a stream of records carrying flow- or packet-based information. The records may be encoded as IPFIX Data Records in any other format.

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.

Original Observation Point

An Observation Point is a location in the network where IP packets are observed, as received by the IPFIX Mediation. Examples include: a (set of) specific exporter(s), a (set of)

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.

Anonymisation Record

A record, defined by the Anonymisation Options Template in section Section 6.1, that defines the properties of the anonymisation applied to a single Information Element within a single Template or Options Template.

Anonymised Data Record

A Data Record within a Data Set containing at least one Information Element with anonymised values. The Information Element(s) within the Template or Options Template describing

Internet-Draft <Protocol for IPFIX Mediations> October 2010
this Data Record SHOULD have a corresponding Anonymisation
Record.

Intermediate Anonymisation Process

An intermediate process which takes Data Records and and
transforms them into Anonymised Data Records.

3. Specifications

This section describes the IPFIX specifications for Mediation:
more specifically, specifications for generic Intermediate
Processes. Possible specific Intermediate Processes are:
Intermediate Conversion Process, Intermediate Aggregation
Process, Intermediate Correlation Process, Intermediate
Selection Process, Intermediate Anonymization Process.

For a specific Intermediate Process, the specifications in the
following reference MUST be followed, on the top of the
specifications in this document:

- For the Intermediate Aggregation Process, the specifications
in [IPFIX-MED-AGGR] MUST be followed.
- For the Intermediate Selection Process, the specifications in
[IPFIX-MED-FLOWSEL] MUST be followed.
EDITOR'S NOTE: actually, there is no MUST/SHOULD/MAY in
[IPFIX-MED-FLOWSEL], which seems to be a list of required
Information Elements.
- For the Intermediate Anonymization Process, the specifications
in [IPFIX-MED-ANON] should be considered as guidelines as
[IPFIX-MED-ANON] is an experimental RFC.

Note that no specific document deals with the Intermediate
Conversion Process at the time of this publication.

These new specifications, which are more specific compared to
[RFC5101], are described with the key words described in
[RFC2119].

3.1. Encoding of IPFIX Message Header

The format of the IPFIX Message Header is shown in Figure A.
Note that the format is similar to the IPFIX Message in
[RFC5101], but some field definitions (for the example, the
Export Time) have been updated in the context of the IPFIX
Mediator.

A 32-bit identifier of the Observation Domain that is locally unique to the Exporting Process. The Exporting Process uses the Observation Domain ID to uniquely identify to the Collecting Process the Observation Domain that metered the Flows. It is RECOMMENDED that this identifier is also unique per IPFIX Device. Collecting Processes SHOULD use the Transport Session and the Observation Domain ID field to separate different export streams originating from the same Exporting Process. The Observation Domain ID SHOULD be 0 when no specific Observation Domain ID is relevant for the entire IPFIX Message. For example, when exporting the Exporting Process Statistics, or in case of hierarchy of Collector when aggregated data records are exported.

Note: the Observation Domain Management is discussed in section 3.4.1.

3.2. Template Management

3.2.1. Template Management Without Template Record Change

The first case is a situation where the IPFIX Mediator, typically an IPFIX Distributor, relays an (Options) Template without changing its content.

As in [RFC5101], the Template IDs are unique per Exporter, per Transport Session, and per Observation Domain. As there is no guarantee that, for similar Template Records, the Template IDs received on the incoming Transport Session and exported to the outgoing Transport Session would be same, the IPFIX Mediator MUST maintain a mapping database between received and exported (Options) Template Records:

- for each Received (Options) Template Record: Template Record Flow Keys and non Flow Keys, Template ID, Original Exporter, Observation Domain, and Transport Session
- for each Exported (Options) Template Record: Template Record Flow Keys and non Flow Keys, Template ID, Collector, Observation Domain, and Transport Session

If an IPFIX Mediator receives an IPFIX Withdrawal Message for a (Options) Template Record that is not used anymore in any outgoing Transport Sessions, the IPFIX Mediator SHOULD export

Internet-Draft <Protocol for IPFIX Mediations> October 2010
the appropriate IPFIX Withdrawal Message(s) on the outgoing
Transport Session, and remove the corresponding entry in its
mapping database.

If a (Options) Template Record is not used anymore in outgoing
Transport Session, it MUST be withdrawn with an IPFIX Template
Withdrawal Message on that specific outgoing Transport Session,
and remove the corresponding entry in its mapping database.

If an incoming Transport Session is gracefully shutdown or
reset, the (Options) Template Records corresponding to that
Transport Session MUST be removed from the mapping database.

3.2.2. Template Management With Template Record Change

The second case is a situation where the IPFIX Mediator,
typically containing an Intermediate Conversion Process,
Intermediate Aggregation Process [IPFIX-MED-AGGR], or
Intermediate Anonymization Process in case of black-marker
anonymisation [IPFIX-MED-ANON], generates new (Options) Template
Records based what it receives from the Original Exporter(s),
and based on the Intermediate Process function.

In such a situation, the IPFIX Mediator doesn't maintain a
mapping database between received and exported (Options)
Template Records, as it generates its own series of (Options)
Template Records.

3.3. Time Management

The IPFIX Message Header "Export Time" field is the time in
seconds since 0000 UTC Jan 1, 1970, at which the IPFIX Message
Header leaves the IPFIX Mediator. However, in the specific case
of an IPFIX Mediator containing an Intermediate Conversion
Process, the IPFIX Mediator MAY keep the export time received
from the incoming Transport Session.

It is RECOMMENDED that Mediators handle time using absolute
timestamps (e.g. flowStartSeconds, flowStartMilliseconds,
flowStartNanoseconds), which are specified relative to the UNIX
epoch (00:00 UTC 1 Jan 1970), where possible, rather than
relative timestamps (e.g. flowStartSysUpTime,
flowStartDeltaMicroseconds), which are specified relative to
protocol structures such as system initialization or message
export time.

The latter are difficult to manage for two reasons. First, they require constant translation, as the system initialization time of an intermediate system and the export time of an intermediate message will change across mediation operations. Further, relative timestamps introduce range problems. For example, when using the flowStartDeltaMicroseconds and flowEndDeltaMicroseconds Information Elements [RFC5102], the Data Record must be exported within a maximum of 71 minutes after its creation. Otherwise, the 32-bit counter would not be sufficient to contain the flow start time offset. Those time constraints might be incompatible with some of the Intermediate Processes: Intermediate Aggregation Process (temporal) and Intermediate Correlation Process, for example.

When an Intermediate Aggregation Process aggregates information from different Flow Records, the typical reporting times SHOULD BE the minimum of the start times and the maximum of the end times. However, if the Flow Records do not overlap, i.e. if there is a time gap between the times in the Flow Records, then the report may be inaccurate. The IPFIX Mediator is only reporting what it knows, on the basis of the information made available to it - and there may not have been any data to observe during the gap. Then again, if there is an overlap in timestamps, there's the potential of double-accounting: different Observation Points may have observed the same traffic simultaneously. Therefore, as there is not a single rule that fits all different situations, the precise rules of applying the Flow Record timestamps in IPFIX Mediators is out of the scope of this document. However, some more specifications related to the specific case of aggregation in space and time are specified in [IPFIX-MED-AGGR], and MUST be followed.

3.4. Observation Point Management

Depending on the use case, top Collectors may need to receive the Original Observation Point(s), otherwise it may wrongly conclude that the IPFIX Device exporting the Flow Records to him, i.e. the IPFIX Mediator, directly observed the packets that generated the Flow Records. Two new Information Element are introduced to solve this use case: originalExporterIPv4Address and originalExporterIPv6Address.

In the IPFIX Mediator, the Observation Point(s) may be represented by:

- A single Original Exporter (represented by the originalExporterIPv4Address or originalExporterIPv6Address Information Elements)
- A list of Original Exporter (represented by the originalExporterIPv4Address or originalExporterIPv6Address Information Elements_)
- A list of Original Exporter (represented by the originalExporterIPv4Address or originalExporterIPv6Address Information Elements), along with the associated interface (represented by the ingressInterface and/or egressInterface)
- A list of Original Exporter (represented by the originalExporterIPv4Address or originalExporterIPv6Address Information Elements), along with the associated line card id (represented by the lineCardId)
- Any combination or list of Information Elements representing Observation Points.

Some Information Elements characterizing the Observation Point may be added. For example, the flowDirection Information Element specifies the direction of the observation, and, as such, characterizes the Observation Point.

Any combination of the above examples is possible. For example, a Original Observation Point composed of:

```
exporterIPv4Address 192.0.2.1
exporterIPv4Address 192.0.2.2,
    interface ethernet 0, direction ingress
    interface ethernet 1, direction ingress
    interface serial 1, direction egress
    interface serial 2, direction egress
exporterIPv4Address 192.0.2.3,
    lineCardId 1, direction ingress
```

If the Original Observation Point composed of a list exported from the IPFIX Mediator, then the IPFIX Structured Data [IPFIX-STRUCT] MUST be used to encode it.

The most generic way to export the Original Observation Point is to use a subTemplateMultiList, with the semantic "exactlyOneOf". Taking back the previous example, the following encoding can be used:

```
Template Record 1: exporterIPv4Address
Template Record 2: exporterIPv4Address, basicList of
    ingressInterface, flowDirecdtion
```

The Original Observation Point is modeled with the Data Records corresponding to either Template Record 1, Template Record 2, or Template Record 3 but not more than one of these ("exactlyOneOf" semantic). This implies that the Flow was observed at exactly one of the Observation Points reported.

When an IPFIX Mediator exports the Original Observation Domain, it SHOULD export other information indicating that an IPFIX Mediator certifies the original exporter IP address. ExporterCertificate in [RFC5655] can be used in that case. And also, another Information Element indicating that certifies that an IPFIX Mediator is required, just like mediatorCertificate.

3.4.1. Observation Domain Management

In terms of Observation Domain management, there are two types of Intermediate Process function. The first one maintains the Observation Domain information, while the second one cannot maintain it. Examples of the second type include mixing Data Records from multiple IPFIX Messages received from multiple Observation Domains, or generating new Data Records from the result of some intermediate function on Data Records from multiple IPFIX Messages received from multiple Observation Domains.

From the two types of Intermediate Process function discussed in this section, a single specification can be deduced, as already specified in 3.1. : The Observation Domain ID SHOULD be 0 when no specific Observation Domain ID is relevant for the entire IPFIX Message.

3.5. Specific Reporting Requirements

Some specific Options Templates and Options Template Records are necessary to provide extra information about the Flow Records and about the Metering Process.

The Options Template Records defined in these subsections, which impose some constraints on the Metering Process and Exporting Process implementations in Intermediate Processes, MAY be

Internet-Draft <Protocol for IPFIX Mediations> October 2010
implemented. If implemented, the specific Option Templates
SHOULD be implemented as specified in these subsections.

The minimum set of Information Elements is always specified in these Specific IPFIX Options Templates. Nevertheless, extra Information Elements may be used in these specific Options Templates.

3.5.1. The Flow Keys Options Template

Exactly like the IPFIX protocol [RFC5101], the Flow Keys Option Template specifies the structure of a Data Record for reporting the Flow Keys of reported Flows. A Flow Keys Data Record extends a particular Template Record that is referenced by its templateId identifier. The Template Record is extended by specifying which of the Information Elements contained in the corresponding Data Records describe Flow properties that serve as Flow Keys of the reported Flow.

The Flow Keys Option Template SHOULD contain the following Information Elements that are defined in [RFC5102]

templateId	An identifier of a Template. This Information Element MUST be defined as a Scope Field.
------------	---

flowKeyIndicator	Bitmap with the positions of the Flow Keys in the Data Records.
------------------	---

When any Intermediate Process changes the Flow Keys, the Flow Keys Option Template MUST include the new set of Flow Keys. Typically, an Intermediate Aggregation Process keeps or reduces the number of Flow Keys

3.5.2. IPFIX Protocol Options Template

The "Metering Process Statistics Options Template", "The Metering Process Reliability Statistics Options Template", and "The Exporting Process Reliability Statistics Options Template", as specified in [RFC5101], SHOULD be implemented on the IPFIX Mediator.

Refer to the document specifying a particular Intermediate Process type for specific values for these Options Template Records. For example, in case of an Intermediate Aggregation Process, [IPFIX-MED-AGGR] must specify which values to insert into the fields of "Metering Process Statistics Options Template", "The Metering Process Reliability Statistics Options

3.5.3. IPFIX Mediator Options Template

There is no need for a specific Options Template for the IPFIX Mediator; instead, each Intermediate Process type requires some particular metadata. For example, a specification of IPFIX flow anonymisation including an Options Template for the export of metadata about anonymised flows is described in [IPFIX-MED-ANON]; when anonymising Flows Records, IPFIX Mediators SHOULD add the Options Template specified therein to annotate the exported dataTransport Session Management Sctp [RFC4960] using the PR-SCTP extension specified in [RFC3758] MUST be implemented by all compliant IPFIX Mediator implementations. UDP [UDP] MAY also be implemented by compliant IPFIX Mediator implementations. TCP [TCP] MAY also be implemented by IPFIX Mediator compliant implementations.

PR-SCTP SHOULD be used in deployments where IPFIX Mediators and Collectors are communicating over links that are susceptible to congestion. PR-SCTP is capable of providing any required degree of reliability.

TCP MAY be used in deployments where IPFIX Mediators and Collectors communicate over links that are susceptible to congestion, but PR-SCTP is preferred due to its ability to limit back pressure on Exporters and its message versus stream orientation.

UDP MAY be used, although it is not a congestion-aware protocol. However, the IPFIX traffic between IPFIX Mediator and Collector MUST run in an environment where IPFIX traffic has been provisioned for, or is contained through some other means.

3.6. The Collecting Process's Side

An IPFIX Mediator MUST produce IPFIX Messages understandable by a RFC5101-compliant IPFIX Collector, with the additional specification in the IPFIX Structured Data [IPFIX-STRUCT].

Therefore the Collecting Process on the top Collector MUST support the IPFIX protocol [RFC5101] and the IPFIX Structured Data [IPFIX-STRUCT].

EDITOR'S NOTE: What about the accuracy of aggregated Flow Records with the sampling rates? With different sampling rates?

In an IPFIX Mediation, aggregation for Flow Records with same sampling rate and same sampling algorithm is recommended. In that case, an IPFIX Mediator can export this sampling rate and sampling algorithm, and other accuracy statistics data, part of the PSAMP Report Interpretation [RFC5476].

In the case where the Mediation aggregates Flow Records with different sampling functions and/or sampling rates, some more research is required to determine the right sampling function and/or sampling rate to export from the IPFIX Mediator. Therefore, this document doesn't describe any specifications, or even guidelines.

4. New Information Elements

- originalExporterIPv4Address and originalExporterIPv6Address
- EDITOR'S NOTE: to be discussed
- originalObservationDomainId?
 - mediatorCertificate?

EDITOR'S NOTE: Maybe the following ones should be defined in a specific flow aggregation draft:

- Maximum counter or minimum counter for packets or bytes
- activeTime and inactiveTime for Flow aggregation

5. Security Considerations

The same security considerations as for the IPFIX Protocol [RFC5101] apply.

As they act as both IPFIX Collecting Processes and Exporting Processes, the Security Considerations for IPFIX [RFC5101] apply as well to Mediators. The Security Considerations for IPFIX Files [RFC5655] apply as well to IPFIX Mediators that write IPFIX Files or use them for internal storage. However, there are a few specific considerations that IPFIX Mediator implementations must take into account in addition.

By design, IPFIX Mediators are "men-in-the-middle": they intercede in the communication between an Original Exporter (or another upstream Mediator) and a downstream Collecting Process. This has two important implications for the level of confidentiality provided across an IPFIX Mediator, and the ability to protect data integrity and Original Exporter authenticity across a Mediator. We address these in the following subsections.

5.1. Avoiding Security Downgrade

An IPFIX Mediator that accepts IPFIX Messages over a Transport Session protected by TLS or DTLS, and which then exports IPFIX Messages derived there from in cleartext, is a potentially serious vulnerability in an IPFIX infrastructure. While this is potentially acceptable in the specific case of an IPFIX Mediator at the border of an administrative domain accepting IPFIX Messages from outside the domain and re-exporting derived information via an internal network protected by other means, in the general case this situation SHOULD be avoided.

Therefore, an IPFIX Mediator that receives IPFIX Messages from an upstream Exporting Process protected using TLS or DTLS MUST provide for sending of IPFIX Messages resulting from the Intermediate Process to a downstream Collecting Process using TLS or DTLS. It MAY allow for the configuration of unprotected export of such IPFIX Messages, but in this case it MUST warn the administrator that the exported IPFIX Messages will not be protected, and that this could result in the leakage of information deemed by the Original Exporter to be worth protecting.

5.2. End-to-End Assertions for Mediators

Because the Transport Session between an IPFIX Mediator and an Original Exporter is independent from the Transport Session between the Mediator and the downstream Collecting Process, there exists no method via TLS to assert the identity of the original Exporting Process downstream. However, an IPFIX Mediator, which modifies the stream of IPFIX Messages sent to it, is by definition a trusted entity in the infrastructure. Therefore, the IPFIX Mediator's signature on an outgoing Transport Session can be treated as an implicit assertion that the Original Exporter was positively identified by the Mediator and that the source information it received was trustworthy.

Internet-Draft <Protocol for IPFIX Mediations> October 2010
However, IPFIX Mediators must in this circumstance take care not to provide an inappropriate upgrade of trust.

Therefore, an IPFIX Mediator SHOULD NOT sign a Transport Session to a downstream Collector unless ALL the Original Exporters from which the information to be exported is derived were positively identified by the Mediator by its certificate. An exception to this case is the reverse of the special case in the previous subsection: an IPFIX Mediator that accepts information from within a trusted domain via an internal network protected by other means MAY use TLS or DTLS to protect the Transport Session to a downstream Collector outside the domain.

[EDITOR OPEN ISSUE: We might want to use exporterCertificate and (optionally) collectorCertificate from [RFC5655] here, but I think they need a new Mediator-specific template if so. If we were to use the templates defined by 5655, it would look like this:

If the X.509 certificates used to protect a Transport Session between an Original Exporter and an IPFIX Mediator are required downstream, an IPFIX Mediator MAY use the exporterCertificate and the collectorCertificate Information Elements with the Export Session Details Options Template defined in Section 8.1.3 of [RFC5655] or the Message Details Options Template defined in Section 8.1.4. of [RFC5655] in order to export this information downstream. However, in this case, the IPFIX Mediator is making an implicit assertion that the upstream Session was properly protected and therefore trustworthy, and as such MUST protect the Transport Session to the downstream Collector using TLS or DTLS, as well.

6. IANA Considerations

This document specifies three new IPFIX Information Elements: the applicationDescription, applicationTag and the applicationName.

New Information Elements to be added to the IPFIX Information Element registry at [IANA-IPFIX] are listed below.

EDITOR'S NOTE: the XML specification in Appendix A must be updated with the elementID values allocated below.

Name: originalExporterIPv4Address

Description:

The IPv4 address used by the Exporting Process on the Original Observation Point. This is used by the Exporting Process (on the Mediation) to identify the Exporter in cases where the identity of the Exporter may have been obscured by the use of a proxy, or in cases where the IPFIX Mediation must export the Original Observation Point to a top Collector.

Abstract Data Type: ipv4Address

Data Type Semantics: identifier

ElementId: XXX

Status: current

6.2. originalExporterIPv6Address

Name: originalExporterIPv6Address

Description:

The IPv6 address used by the Exporting Process on the Original Observation Point. This is used by the Exporting Process (on the Mediation) to identify the Exporter in cases where the identity of the Exporter may have been obscured by the use of a proxy, or in cases where the IPFIX Mediation must export the Original Observation Point to a top Collector.

Abstract Data Type: ipv6Address

Data Type Semantics: identifier

ElementId: YYY

Status: current

EDITOR'S NOTE: maybe some more IEs

7. References

7.1. Normative References

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

[RFC3758] Stewart, R., Ramalho, M, Xie, Q., Tuexen, M., and P. Conrad, "Stream Control Transmission Protocol (SCTP), Partial Reliability Extension", May 2004

[RFC4960] Stewart, R., Ed., "Stream Control Transmission Protocol", RFC 4960, September 2007.

- Internet-Draft <Protocol for IPFIX Mediations> October 2010
- [RFC5101] Claise, B., Ed., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", RFC 5101, January 2008.
- [RFC5102] Quittek, J., Bryant, S., Claise, B., Aitken, P., and J. Meyer, "Information Model for IP Flow Information Export", RFC 5102, January 2008.
- [RFC5655] Trammell, B., Boschi, E., Mark, L., Zseby, T., and A. Wagner, "Specification of the IP Flow Information Export (IPFIX) File Format", RFC 5655, October 2009.
- [IPFIX-MED-FLOWSEL] Peluso, L., Zseby, T., D'antonio, S., and M. Molina, "Flow Selection Techniques", draft-ietf-ipfix-flow-selection-tech-02.txt, Internet-Draft work in progress, June 2010.
- [IPFIX-MED-AGGR] Trammell, B., Boschi, E., and A. Wagner, "Exporting Aggregated Flow Data using the IP Flow Information Export (IPFIX) Protocol", draft-trammell-ipfix-a9n-01.txt, Internet-Draft work in progress, October 2010.
- [IPFIX-STRUCT] Claise, B., Dhandapani, G., Aitken, P., and S. Yates, "Export of Structured Data in IPFIX", draft-ietf-ipfix-structured-data-03.txt, Internet-Draft work in progress, October 2010.

7.2. Informative References

- [TCP] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [UDP] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.
- [RFC3917] Quittek, J., Zseby, T., Claise, B., and S. Zander, "Requirements for IP Flow Information Export", RFC 3917, October 2004
- [RFC3954] Claise, B. (Ed), "Cisco Systems NetFlow Services Export Version 9", RFC 3954, October 2004

Internet-Draft <Protocol for IPFIX Mediations> October 2010
[RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture Model for IP Flow Information Export", RFC5470, March 2009

[RFC5472] Zseby, T., Boschi, E., Brownlee, N., and B. Claise, "IP Flow Information Export (IPFIX) Applicability", RFC 5472, March 2009

[RFC5476] Claise, B., Quittek, J., and A. Johnson, "Packet Sampling (PSAMP) Protocol Specifications", RFC 5476, March 2009.
[RFC5982] Kobayashi, A. (Ed), Claise, B. (Ed), "P Flow Information Export (IPFIX) Mediation: Problem Statement", RFC 5982, August 2010.

[IPFIX-MED-FMWK] Kobayashi, A., Claise, B., and K. Ishibashi, "IPFIX Mediation: Framework", draft-ietf-ipfix-mediators-framework-09, Internet-Draft work in progress, October 2010.

[IPFIX-MED-ANON] Boschi, E., Trammell, B. "IP Flow Anonymisation Support", draft-ietf-ipfix-anon-05.txt, Internet-Draft work in progress, October 2010.

[IANA-IPFIX] <http://www.iana.org/assignments/ipfix/ipfix.xhtml>

8. Author's Addresses

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
Diegem 1813
Belgium

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

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

Brian Trammell
Hitachi Europe
c/o ETH Zurich
Gloriastrasse 35
8092 Zurich
Switzerland

Phone: +41 44 632 70 13
EMail: brian.trammell@hitachi-eu.com

9. Appendix A. Additions to XML Specification of IPFIX Information Elements

This appendix contains additions to the machine-readable description of the IPFIX information model coded in XML in Appendix A and Appendix B in [RFC5102]. Note that this appendix is of informational nature, while the text in Section 6. (generated from this appendix) is normative.

The following field definitions are appended to the IPFIX information model in Appendix A of [RFC5102].

```
<field name=" originalExporterIPv4Address "  
      dataType="identifier "  
      group="config"  
      elementId="XXX" applicability="all" status="current">  
  <description>  
    <paragraph>  
      The IPv4 address used by the Exporting Process on the  
      Original Observation Point. This is used by the  
      Exporting Process (on the Mediation) to identify the  
      Exporter in cases where the identity of the Exporter  
      may have been obscured by the use of a proxy, or in  
      cases where the IPFIX Mediation must export the  
      Original Observation Point to a top Collector.  
    </paragraph>  
  </description>  
</field>
```

```
<field name=" originalExporterIPv6Address "  
      dataType="identifier "  
      group="config"  
      elementId="XXX" applicability="all" status="current">
```

```
<description>
  <paragraph>
    The IPv6 address used by the Exporting Process on the
    Original Observation Point. This is used by the
    Exporting Process (on the Mediation) to identify the
    Exporter in cases where the identity of the Exporter
    may have been obscured by the use of a proxy, or in
    cases where the IPFIX Mediation must export the
    Original Observation Point to a top Collector.
  </paragraph>
</description>
</field>
```