IPFIX Working Group
Internet-Draft
Intended status: Informational
Expires: August 9, 2009

A. Kobayashi, ED.
NTT PF Lab.

February 5, 2009

IPFIX Mediation: Problem Statement
draft-ietf-ipfix-mediators-problem-statement-02

Status of this Memo

Copyright Notice

Abstract

Flow-based measurement is a popular method for various network monitoring usages. The sharing of flow-based information for monitoring applications having different requirements raises some open issues in terms of scalability, reliability, and flexibility that IPFIX Mediation may help resolve. IPFIX Mediation covers two classes of mediation: context mediation for traffic data and transport mediation for transport protocols. This document describes the problems that network administrators have been facing and the applicability of IPFIX Mediation.

Table of Contents

1. Introduction

While the IPFIX requirements defined in [RFC3917] mention an intermediate function, such as an IPFIX Proxy or an Concentrator, there is no document to define the function called IPFIX Mediation. IPFIX Mediation is a generic function that covers context mediation for traffic data and transport mediation for IPFIX transport protocols that do not affect content. We describes the general problems that network administrators have been facing and several applicable IPFIX Mediation categories along with specific terminology (IPFIX Proxy, Concentrator, etc.). Furthermore, we describe the problems of IPFIX Mediation with regard to implementation. These problems can be solved by making additional specifications that do not affect the present IPFIX protocol specifications defined in [RFC5101].

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 introduces general problems related to flow-based measurement. Section 5 describes some applicable examples where IPFIX Mediation would benefit from solutions to such problems. Finally, section 6 describes the problems an implementation of an IPFIX Mediation device might face.

2. Terminology and Definition

The terms in this section are in line with those in the IPFIX Protocol specifications [RFC5101] and the PSAMP specification document [I-D.ietf-psamp-protocol]. The terms Observation Point, Observation Domain, Flow Key, Flow Record, Exporting Process, Exporter, IPFIX Device, Collecting Process, Collector, IPFIX Message, Metering Process, and Information Element are defined in the IPFIX protocol specifications [RFC5101], while the term Packet Report is defined in the PSAMP specification document [I-D.ietf-psamp-protocol]. Additional terms required for the IPFIX Mediation are also defined here. All these terms have an initial capital letter in this document.

IPFIX Mediation

IPFIX Mediation is a function that can be applied to individual Data Records and/or Template Records or to entire IPFIX Messages. IPFIX Mediation offers one or multiple capabilities.

* content mediation that changes Flow information
  + aggregating Data Records based on a new set of Flow Key fields
  + correlating a set of Data Records for creating new metrics
  + filtering and selecting Data Records
  + modifying Data Records and/or Template Records, which includes these functions:
    - changing the value of specified Information Elements
    - adding new Information Elements by deriving further Flow or packet properties from existing fields or calculating new metrics
    - deleting specified Information Elements
* transport mediation that does not affect content
  + changing the transport protocol that carries IPFIX Messages
  + rerouting entire IPFIX Messages to an appropriate Collecting Process

+ replicating Data Records and Template Records or entire IPFIX Messages

IPFIX Mediation can be included in any IPFIX Devices, such as routers, switches, and network management systems (NMS).

IPFIX Mediator

An IPFIX Mediator is an IPFIX Device that contains one or more IPFIX Mediation capabilities.

Original Exporter

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

IPFIX Proxy

An IPFIX Proxy is an IPFIX Mediation that relays incoming Transport Sessions to one or multiple Collectors. The protocols used at the input and the output may be different, which implies that IPFIX Messages, Data Records, or Template Records need to be encoded, e.g., converting legacy protocol into IPFIX.

IPFIX Concentrator

An IPFIX Concentrator is an IPFIX Mediation that receives Flow Records/Packet Reports, aggregates them, then exports the aggregated Flow Records.

IPFIX Distributor

An IPFIX Distributor is an IPFIX Mediation that distributes incoming IPFIX Data Records to one or multiple IPFIX Collectors. The decision as to which IPFIX Collector a Data Record is exported can be determined by filtering certain field values or other properties derived from the Data Record.

IPFIX Masquerading Proxy

An IPFIX Masquerading Proxy is an IPFIX Mediation that screens out parts of input Data Records according to configured policies. It can thus, for example, hide the network topology information or customers' IP addresses.

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 out of an IPFIX Exporting Process to a Collecting Process is defined in [I-D.ietf-ipfix-architecture], 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]. [I-D.ietf-ipfix-mib] specifies the IPFIX Management Information Base. Finally, [I-D.ietf-ipfix-as] 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 [I-D.ietf-ipfix-file].

3.2. PSAMP Documents Overview

The framework for packet selection and reporting [I-D.ietf-psamp-framework] 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, filtering, and hashing) standardized by PSAMP are described in [I-D.ietf-psamp-sample-tech]. The PSAMP protocol [I-D.ietf-psamp-protocol] 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 [I-D.ietf-psamp-info]. [I-D.ietf-psamp-mib] describes the PSAMP Management Information Base.

4. Problem Statement

Network administrators generally face the problems of flow-based measurement for scalability, reliability, and flexibility, and some techniques, such as sampling, aggregating and replicating, have already been developed. The problems consist of optimizing the resources of the measurement system while pursuing appropriate conditions, such as data accuracy, flow granularity, and reliability. The conditions depend on two factors.

- capacity of measurement system
  This consists of the bandwidth of the management network, the storage capacity, and the performances of the collecting devices and exporting devices.

- requirement for given applications
  This depends on the purpose of the application, such as traffic engineering, detecting anomaly traffic, and accounting.

The recent continued IP traffic growth has been overwhelming the capacity of measurement system, and multi-purposing applications and the heterogeneous environment have further contributed to a complex situation. The following sub-sections explain problems related to these two factors.

4.1. Approach for IP Traffic Growth

Enterprise or service provider networks already have multiple 10 Gb/s links, their total traffic exceeding 100 Gb/s. In the near future, broadband users' traffic will increase by approximately 40% every year according to [TRAFGRW]. When operators monitor traffic of 500 Gb/s with a sampling rate of 1/1000, the amount of exported Flow Records from Exporters could exceed 50 kFlows/s. This value is beyond the ability of a single Collector.

To deal with this problem, traffic data reduction techniques, such as sampling or aggregating, have been generally implemented in exporting devices. These techniques lead to coarse flow granularity or low data accuracy, resulting in Flows with small traffic volumes that could easily get lost. Administrators would no longer be able to investigate traffic change and anomaly traffic, both of which can currently be detected, unless the collecting infrastructure is improved.

This implies the necessity of a large-scale collecting infrastructure and other traffic data reduction techniques other than packet-based sampling and selection techniques.

4.2. Approach to Multifaceted Traffic Measurement

A set of conditions (flow granularity and data accuracy) may meet the requirements of some applications, such as traffic engineering, but would not meet the requirements of other applications, such as accounting and QoS performance. Therefore, with a single set of conditions, multifaceted traffic measurement cannot be accomplished.

To cope with the issue, a exporting device needs to export traffic data with strictest condition (fine flow granularity and high data accuracy) required by one of applications. However, it brings about increasing the load on a exporting device and a collecting device.

4.3. Approach to Heterogeneous Environment

Network administrators use exporting devices from various vendors and of various software versions or device type (router, switch, or probe) in a single network domain. This heterogeneous environment leads to differences in capability, performance, and data format. For example, a probe and a switch cannot retrieve packet property information from a route table.

To deal with this problem, a collecting device needs to absorb the differences. However, equipping all collecting devices with this extra function is difficult. A sophisticated solution that introduces individual modules separate from specific devices is necessary.

4.4. Summary

In optimizing the resources of a measurement system, it is important to use traffic data reduction techniques at the possible initial phase, e.g., exporting devices, of the whole system. However, this implementation is made difficult by heterogeneous environment of exporting devices.

It implies that the exporter-collector structure model has limitations, and a mediation function as another functional block is necessary. The next section shows the limitation of the exporter-collector structure model and the benefit of IPFIX Mediation according in applicable examples.

5. Applicable Examples

5.1. Adjusting Flow Granularity

The simplest types of Flows are those comprised of packets all having a fixed IP-quintuple of protocol, source and destination IP addresses, and source and destination port numbers. However, a shorter Flow Key, such as a triple, a double, or a single Flow Key, such as a network prefix, peering AS number, or BGP Next-Hop, creates more aggregated Flow Records. This is especially useful for measuring traffic exchange in an entire network domain and for easily adjusting the performance of a Collector.

Implementation analysis:

Implementations for this case depend on where Flow granularity is adjusted. More suitable implementations uses the configurable Metering Process in Original Exporters. The cache in the Metering Process can specify its own set of Flow Keys and extra fields. The Original Exporter thus creates directly aggregated Flow Records.

In the case where an unconfigurable Metering Process creating IP-quintuple Flow Records exists in a line interface module, IPFIX Mediation in another module can be applied between the Metering Process and an Exporting Process.

In the case where an Original Exporter creating IP-quintuple Flow Records exists, an IPFIX Concentrator can be applied between the Original Exporter and an IPFIX Collector.

5.2. Hierarchical Collecting Infrastructure

As an approach to large-scale measurement systems, a hierarchical structure is useful for increasing the capacity.

Implementation analysis:

One possible implementation for this case uses an IPFIX Concentrator. An IPFIX Concentrator with storage capability also makes a most useful distributed-collection system.

5.3. Correlation of Data Records

The correlation of Data Records provides new metrics, including the following.

- One way delay from the correlation of Packet Reports from different Exporters on the path.

- Rate-limiting ratio from the correlation of Data Records with the same Flow Key observed at incoming/outgoing interfaces.

- Average/maximum/minimum values from correlating multiple Data Records.

Implementation analysis:

One possible implementation for this case uses an IPFIX Mediation located between the Metering Processes and Exporting Processes or between the Original Exporters and IPFIX Collectors.

5.4. Time Composition

Time composition is defined as the aggregation of consecutive Data Records with identical Flow Key values. It leads to the same output as setting a longer active interval timer on Original Exporters. An advantage is that creating new metrics (average, maximum and minimum values) from Flow Records with a shorter interval time enables administrators to keep track of changes that might have happened during the time interval.

Implementation analysis:

One possible implementation for this case uses an IPFIX Mediation located between the Metering Processes and Exporting Processes or between the Original Exporters and IPFIX Collectors.

5.5. Spatial Composition

Spatial composition is defined as the aggregation of Data Records in a set of Observation Points with an Observation Domain, across multiple Observation Domains from a single Exporter, or even across multiple Exporters. It is divided into three types.

- Spatial Composition within one Observation Domain

  For example, in the case where a link aggregation exists, Data Records observed at physical interfaces belonging to a same trunk can be merged.

- Spatial Composition across Observation Domains, but within a single Exporter

  For example, in the case where a link aggregation exists, Data

Records observed at physical interfaces belonging to a same trunk grouping beyond the line interface module can be merged.

- Spatial Composition across Exporters

  Data Records observed at different domains, such as the west area and east area of an ISP network, can be merged.

Implementation analysis:

One possible implementation for this case uses an IPFIX Mediation located between the Metering Processes and Exporting Processes or between the IPFIX Exporters and IPFIX Collectors.

5.6. Data Retention

Data retention refers to the storage of traffic data by service providers and commercial organizations. In accordance with European Commission directives, operators are required to retain both IP and voice traffic data, in wired and wireless networks, generated by end users while using a service provider's services. The goal of data retention is to ensure that call detail records and Flow Records are available for the detection, investigation, and prosecution of serious crimes, if necessary. The European Commission directives define the following data retention services:

- Fixed telephony (includes fixed voice calls, voicemail, and conference and data calls)
- Mobile telephony (includes mobile voice calls, voicemail, conference and data calls, SMS, and MMS)
- Internet telephony (includes every multimedia session associated with IP multimedia services)
- Internet e-mail
- Internet access

Data retention for Internet access services in particular requires a measurement system with reliability and huge storage.

Implementation analysis:

Regarding reliability, the most suitable implementation uses the SCTP transport protocol between the Original Exporter and Collector. Otherwise, an IPFIX Proxy next to a legacy exporting device exports traffic data to the final IPFIX Collector through

SCTP.

Regarding huge storage, one possible implemantation uses a decentralized collecting device. If operators need to retrieve specific traffic data, these collecting devices would need to be equipped with IPFIX Mediation capabilities.

[Editor Note]

The authors need to find the data retention reference.

5.7. IPFIX Export from Branch Office

Generally, in large enterprise networks, traffic data from branch offices are gathered in a central office. However, in the long distance branch office case, the bandwidth for transport IPFIX is limited.

Implementation analysis:

One possible implementation for this case uses an IPFIX Concentrator located in a branch office. The IPFIX Concentrator then exports aggregated Flow Records to cope with the bandwidth limitation.

5.8. Distributing Data Records

Recently, several networks have shifted towards integrated networks, such as the pure IP and MPLS, which includes IPv4, IPv6, and VPN traffic. Data Record types (IPv4, IPv6, MPLS, and VPN) need to be analyzed separately and from different perspectives. However, handling them separately without improving the capability of the Collector is difficult. Data Records distributed based on the type can be exported to an appropriate Collector with a specific application, and this results in the distribution of the load among multiple Collectors.

Implementation analysis:

One possible implementation for this case uses the replications of the IPFIX Message in an IPFIX Exporter for multiple IPFIX Collectors. Each Collector then extracts the Data Record required by its own applications. However, this increases the load of the Exporting Process.

A more sophisticated implementation uses an IPFIX Distributor located between the Metering Processes and Exporting Processes or between the Original Exporters and IPFIX Collectors. For example,

in the case of distributing a specific customer's Data Records, an IPFIX Distributor needs to identify the customer networks. The Route Distinguisher (RD), ingress interface, peering AS number, or BGP Next-Hop, or simply the network prefix may be evaluated to distinguish different customer networks. In the following figure, the IPFIX Distributor reroutes Data Records on the basis of the RD value. This system enables each customer's traffic to be inspected independently.

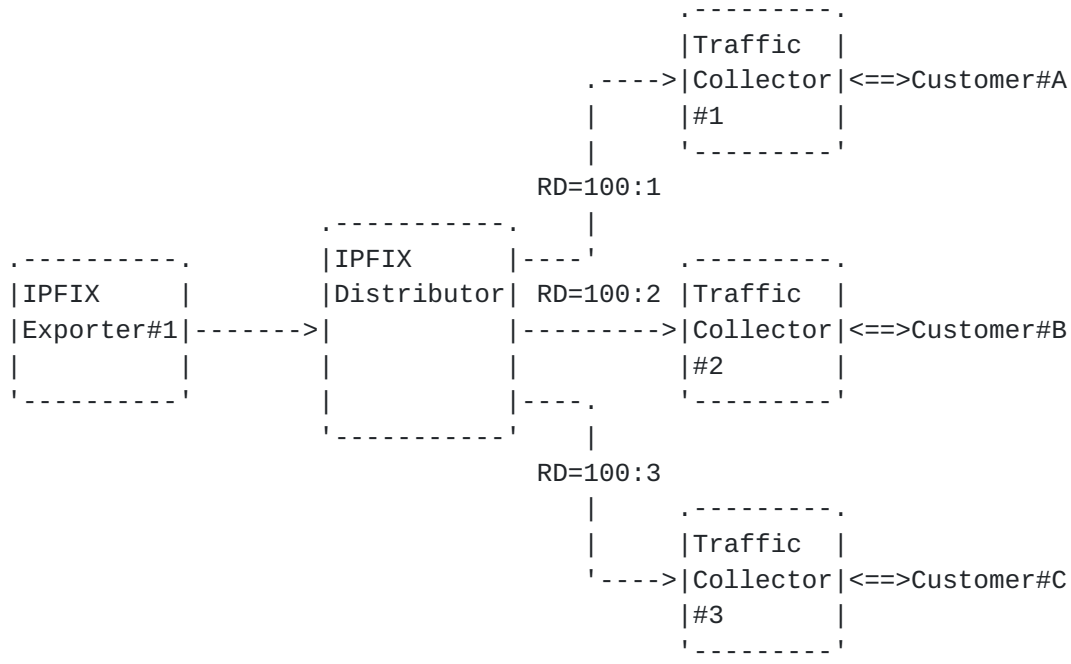

Figure A: Distributing Data Records to Collectors using IPFIX Distributor

5.9. IPFIX Export Across Domains

IPFIX exports across administrative domains can be used to measure traffic for wide-area traffic engineering or to analyze Internet traffic trends. In such cases, administrators need to adhere to privacy protection policies and prevent access to confidential traffic measurements by other people. Typically, anonymization techniques enables the provision of traffic data to other people without violating these policies.

Generally, anonymization modifies a data set to protect the identity of the people or entities described by the data set from being disclosed. It also attempts to preserve sets of network traffic properties useful for a given analysis while ensuring the data cannot be traced back to the specific networks, hosts, or users generating the traffic. For example, IP address anonymization is particularly

important for avoiding the identification of the users, hosts, and routers in a network domain. The details of these anonymization techniques are out of the scope of this document.

Implementation analysis:

One possible implementation for this case uses an anonymization function at the Original Exporter. However, this increases the load of the Metering Process at the Original Exporter. A more flexible implementation uses an IPFIX Masquerading Proxy between the Original Exporter and Collector.

5.10. Flow-based Sampling and Selection

Generally, the distribution of the number of packets per Flow seems to be heavy-tailed. Most types of Flow Records are likely to be small Flows consisting of a small number of packets. The measurement system is overwhelmed with a huge amount of these small Flows. If statistics information of small Flows is exported as merged data by applying a policy or threshold, the load on the measurement system is reduced. Furthermore, if the flow distribution is known, exporting only a subset of the Data Records might be sufficient.

Implementation analysis:

One possible implementation for this case uses an IPFIX Mediation located between the Metering Processes and Exporting Processes or between the Original Exporters and IPFIX Collectors.

5.11. Interoperability between Legacy Protocols and IPFIX

During the migration process from a legacy protocol such as NetFlow [RFC3954] to IPFIX, both NetFlow exporting devices and IPFIX Exporters are likely to coexist in the same network. Operators need to continue measuring the traffic data from legacy exporting devices, even after introducing IPFIX Collectors.

Implementation analysis:

One possible implementation for this case uses an IPFIX Proxy that converts a legacy protocol to IPFIX.

6. Problems with using IPFIX Mediators

In this section, we focus on the problems related to the use of IPFIX Mediators in consideration of implementation.

6.1. Loss of Original Exporter Information

Both the Exporter IP address indicated by the source IP address of the IPFIX session and the Observation Domain ID included in the IPFIX Message header are likely to be lost by an IPFIX Mediator such as IPFIX Concentrator. In some cases, the IPFIX Masquerading Proxy might be made to drop the information. However, in other cases, the information is necessary for indicating the Observation Point information from the viewpoint of the entire network domain. Such information might be necessary for guaranteeing the continuity of the work of the top level Collector. Even if an IPFIX Mediator could, with some new mechanism, notify Collectors of this Observation Point information, older Collectors might not accept it. These Collectors would then wrongly assume that the IP address of the IPFIX Mediator is that of the Original Exporter. The Collector, however, sometimes needs to recognize the Original Exporter (and potentially the Observation Domain and Observation Point as well) whether Data Records go through an IPFIX Mediator or not.

In the following figure, a Collector can identify two IP addresses: 10.1.1.3 (IPFIX Mediator) and 10.1.1.2 (Exporter#2). respectively. The Collector, however, needs to somehow recognize both Exporter#1 and Exporter#2, which are the Original Exporters. Defined notification methods that can be interpreted by Collectors and Mediators are thus necessary.

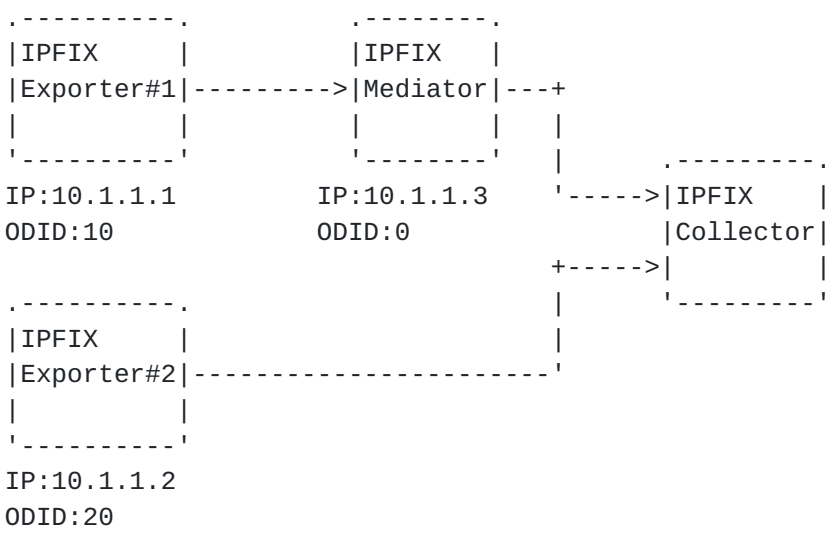

Figure B: Loss of Original Exporter Information.

6.2. Loss of Base Time Information

The Export Time field included in the IPFIX Message header indicates the base time for Data Records. IPFIX Information Elements, described in [RFC5102], have delta time fields that indicate the time difference from the value of the Export Time field. If the Data Records include any delta time fields and the IPFIX Mediator overwrites the Export Time field when sending IPFIX Messages, the delta time fields become meaningless and, because Collectors cannot recognize this situation, wrong time values are propagated.

6.3. Loss of Option Template Information

In some cases, depending on the implementation of the IPFIX Mediators, the information that is reported by the Option Templates could also be lost. If, for example, the sampling rate is not communicated to the Collectors, a Collector would miscalculate the traffic volume. This might lead to crucial problems. Even if an IPFIX Mediator was to simply relay received Option Template Information, the values of its scope fields could become meaningless in the context of a different session. The minimal information to be communicated by an IPFIX Mediator needs to be defined.

6.4. Observation Domain ID and Template ID Management

The Observation Domain ID is locally unique to an Exporting Process, just as the Template ID is unique on the basis of the Transport Session and Observation Domain ID. If IPFIX Mediators were not able to manage the relations among these identifiers and the incoming Transport Session information, and if the Template ID was used in the scope field of Options, the Mediators would, for example, relay wrong values for the scope field and for "Template Withdraw Message". The Collector would thus not be able to interpret the Template ID of "Template Withdraw Message" and of the scope fields of Options. The Collector would then shut down the IPFIX Transport Session.

6.5. Transport Sessions Management

Maintaining relationships between the incoming Transport Sessions and the outgoing ones depends on the Mediator's implementation. If multiple incoming Transport Sessions are relayed to a single outgoing Transport Session, and if the IPFIX Mediators shuts down its outgoing Transport Session, Data Records on other incoming Transport Sessions would not be relayed at all. In the case of resetting of an incoming session, the behavior of the IPFIX Mediator needs to be defined.

For example, an IPFIX Distributor must maintain the state of the incoming Transport Sessions in order to manage the Template ID on its

outgoing Transport Session correctly. In the following figure, even if the Transport Session from Exporter#1 re-initializes, the IPFIX Distributor must maintain the validity of the Template IDs to avoid overlapping the existing ones on the outgoing Transport Session.

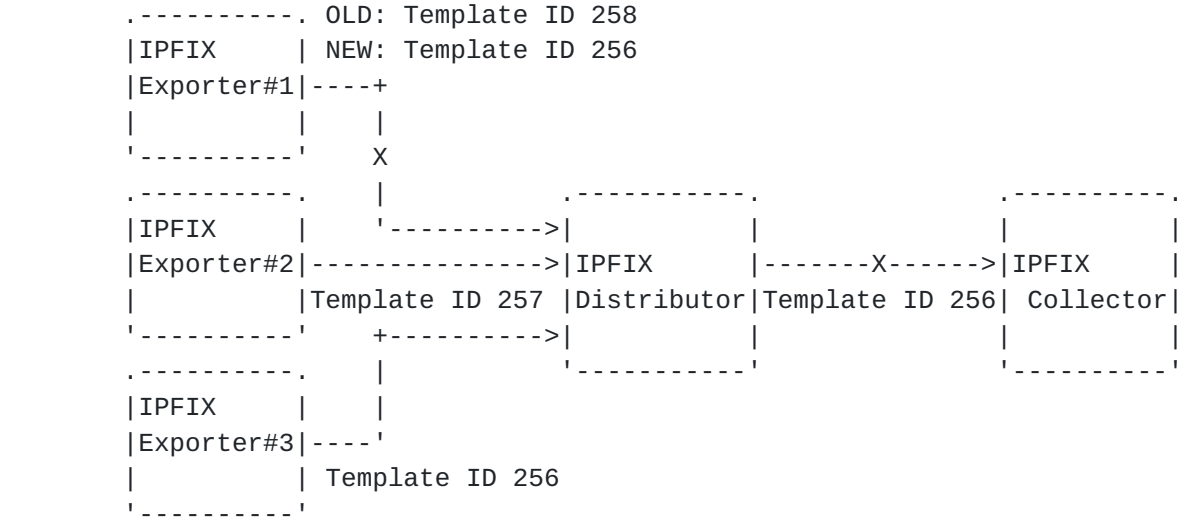

Figure C: Relaying from Multiple Transport Sessions to Single Transport Session.

7. Summary and Conclusion

This document described the problems that network administrators have been facing, the applicability of IPFIX Mediation to these problems, and the problems related to the implementation of IPFIX Mediators. To assist the operations of the Exporters and Collectors, there are various IPFIX Mediation functions from which the administrators may select. Examples of the applicability of IPFIX Mediation are as follows.

- Regarding large-scale measurement system, IPFIX Concentrators or IPFIX Distributors help to achieve traffic analysis with high data accuracy and fine flow granularity even as IP traffic grows. As IPFIX Mediation capabilities, Flow selection sampling, aggregation, and composition are effective.

- Regarding data retention, IPFIX Mediators enhance the reliability, and the storage of the measurement system.

- Regarding the distribution of Data Records, this could be introduced in integrated networks, which mix MPLS VPN and IPv4/ IPv6, more frequently. More sophisticated implementation methods would enhance the distribution's effectiveness.

- Regarding IPFIX Exporting across domains, IPFIX Masquerading Proxies help administrators to anonymize or filter Flow Records/ Packet Reports, preventing privacy violations.

- Regarding interoperability, IPFIX Proxies provide interoperability between legacy protocols and IPFIX, even during the migration period to IPFIX.

As a result, the benefits of IPFIX Mediation become apparent. However, there are still some open issues with the use of IPFIX Mediators.

- Both Observation Point and IPFIX Message header information, such as the Exporter IP address, Observation Domain ID, and Export Time field, might be lost. This data should therefore be communicated between the Original Exporter and Collector via the IPFIX Mediator.

- Data advertised by Option Templates from the Original Exporter, such as the sampling rate and sampling algorithm used, might be lost. If a Collector is not informed of current sampling rates, traffic information might become worthless.

- IPFIX Mediators are required to manage Transport Sessions, Template IDs, and Observation Domain IDs. Otherwise, anomalous IPFIX messages could be created.

These problems stem from the fact that no standards regarding IPFIX Mediation have been set. In particular, the minimum set of information that should be communicated between Original Exporters and Collectors, the interworking between different IPFIX Transport Sessions, and the internal components of IPFIX Mediators should be standardized.

8. Security Considerations

A flow-based measurement system must prevent potential security threats: the disclosure of confidential traffic data, injection of incorrect data, and unauthorized access to traffic data. These security threats of the IPFIX protocol are covered by the security considerations section in [RFC5101] and are true of IPFIX Mediation as well.

And a measurement system must also prevent following security threats related to IPFIX Mediation.

- attacks against IPFIX Mediator

  IPFIX Mediators would be considered a prime target for attacks instead of IPFIX Exporters and Collectors. IPFIX Proxies or Masquerading Proxies need to prevent unauthorized access or denial-of-service (DoS) attacks from untrusted public networks.

- man-in-the-middle attack by untrusted IPFIX Mediator

  The Collector-Mediator-Exporter structure model would increase the risk of the man-in-the-middle attack.

- configuration on 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.

9. IANA Considerations

This document has no actions for IANA.

10. References

10.1. Normative References

[I-D.ietf-ipfix-architecture]
Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", draft-ietf-ipfix-architecture-12 (work in progress) , September 2006.

[I-D.ietf-ipfix-as]
Zseby, T., Boschi, E., Brownlee, N., and B. Claise, "IPFIX Applicability", draft-ietf-ipfix-as-12 (work in progress) , June 2007.

[I-D.ietf-ipfix-file]
Trammell, B., Boschi, E., Mark, L., Zseby, T., and A. Wagner, "Specification of the IPFIX File Format", draft-ietf-ipfix-file-03 (work in progress) , October 2008.

[I-D.ietf-ipfix-mib]
Dietz, T., Claise, B., and A. Kobayashi, "Definitions of Managed Objects for IP Flow Information Export", draft-ietf-ipfix-mib-05 (work in progress) , November 2008.

[I-D.ietf-psamp-framework]
Duffield, N., "A Framework for Packet Selection and Reporting", draft-ietf-psamp-framework-13 (work in progress) , June 2008.

[I-D.ietf-psamp-info]
Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", draft-ietf-psamp-info-11 (work in progress) , October 2008.

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

[I-D.ietf-psamp-protocol]
Claise, B., "Packet Sampling (PSAMP) Protocol Specifications", draft-ietf-psamp-protocol-09 (work in progress) , December 2007.

[I-D.ietf-psamp-sample-tech]
Zseby, T., Molina, M., Duffield, N., Niccolini, S., and F. Raspall, "Sampling and Filtering Techniques for IP Packet Selection", draft-ietf-psamp-sample-tech-11 (work in progress) , July 2008.

[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.

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

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

10.2. Informative References

[TRAFGRW] Cho, K., Fukuda, K., Esaki, H., and A. Kato, "The Impact and Implications of the Growth in Residential User-to-User Traffic", SIGCOMM2006, pp. 207-218, Pisa, Italy, September 2006. .

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
URI: http://www3.plala.or.jp/akoba/

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

Phone: +81-422-59-3978
Email: nishida.haruhiko@lab.ntt.co.jp

Christoph Sommer
University of Erlangen-Nuremberg
Department of Computer Science 7
Martensstr. 3
Erlangen 91058
Germany

Phone: +49 9131 85-27993
Email: christoph.sommer@informatik.uni-erlangen.de
URI: http://www7.informatik.uni-erlangen.de/~sommer/

Falko Dressler
University of Erlangen-Nuremberg
Department of Computer Science 7
Martensstr. 3
Erlangen 91058
Germany

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

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

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

Stephan Emile
France Telecom
2 avenue Pierre Marzin
Lannion, F-22307

Fax: +33 2 96 05 18 52
Email: emile.stephan@orange-ftgroup.com