Internet Engineering Task Force (IETF)
Request for Comments: 7014
Category: Standards Track
ISSN: 2070-1721

S. D'Antonio
Univ. of Napoli "Parthenope"
T. Zseby
CAIDA/FhG FOKUS
C. Henke
Tektronix Communications Berlin
L. Peluso
University of Napoli

September 2013

Flow Selection Techniques

Abstract

The Intermediate Flow Selection Process is the process of selecting a subset of Flows from all observed Flows. The Intermediate Flow Selection Process may be located at an IP Flow Information Export (IPFIX) Exporter or Collector, or within an IPFIX Mediator. It reduces the effort of post-processing Flow data and transferring Flow Records. This document describes motivations for using the Intermediate Flow Selection process and presents Intermediate Flow Selection techniques. It provides an information model for configuring Intermediate Flow Selection Process techniques and discusses what information about an Intermediate Flow Selection Process should be exported.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7014.

Table of Contents

1. Introduction

This document describes Intermediate Flow Selection Process techniques for network traffic measurements. A Flow is defined as a set of packets with common properties, as described in [RFC7011]. An Intermediate Flow Selection Process can be executed to limit the resource demands for capturing, storing, exporting, and post-processing Flow Records. It also can be used to select a particular set of Flows that are of interest to a specific application. This document provides a categorization of Intermediate Flow Selection Process techniques and describes configuration and reporting parameters for them.

This document also addresses configuration and reporting parameters for Flow-state dependent packet selection as described in [RFC5475], although this technique is categorized as packet selection. The reason is that Flow-state dependent packet selection techniques often aim at the reduction of resources for Flow capturing and Flow processing. Furthermore, these techniques were only briefly discussed in [RFC5475]. Therefore, configuration and reporting considerations for Flow-state dependent packet selection techniques have been included in this document.

1.1. Requirements Language

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].

2. Terminology

This document is consistent with the terminology introduced in [RFC7011], [RFC5470], [RFC5475], and [RFC3917]. As in [RFC7011] and [RFC5476], the first letter of each IPFIX specific and Packet Sampling (PSAMP) specific term is capitalized, along with the Intermediate Flow Selection Process specific terms defined here.

* Packet Classification

  Packet Classification is a process by which packets are mapped to specific Flow Records, based on packet properties or external properties (e.g., interface). The properties (e.g., header information, packet content, Autonomous System (AS) number) make up the Flow Key. If a Flow Record for a specific Flow Key value already exists, the Flow Record is updated; otherwise, a new Flow Record is created.

* Intermediate Flow Selection Process

  An Intermediate Flow Selection Process is an Intermediate Process, as defined in [RFC6183] that takes Flow Records as its input and selects a subset of this set as its output. The Intermediate Flow Selection Process is a more general concept than the Intermediate Selection Process as defined in [RFC6183]. While an Intermediate Selection Process selects Flow Records from a sequence based upon criteria-evaluated Flow Record values and only passes on those Flow Records that match the criteria, an Intermediate Flow Selection Process selects Flow Records using selection criteria applicable to a larger set of Flow characteristics and information.

* Flow Cache

  A Flow Cache is the set of Flow Records.

* Flow Selection State

  An Intermediate Flow Selection Process maintains state information for use by the Flow Selector. At a given time, the Flow Selection State may depend on Flows and packets observed at and before that time, as well as other variables. Examples include:

  (i) sequence number of packets and Flow Records;

  (ii) number of selected Flows;

  (iii) number of observed Flows;

  (iv) current Flow Cache occupancy;

  (v) Flow specific counters, lower and upper bounds;

  (vi) Intermediate Flow Selection Process timeout intervals.

* Flow Selector

  A Flow Selector defines the action of an Intermediate Flow Selection Process on a single Flow of its input. The Flow Selector can make use of the following information in order to establish whether or not a Flow has to be selected:

  (i) the content of the Flow Record;

  (ii) any state information related to the Metering Process or Exporting Process;

  (iii) any Flow Selection State that may be maintained by the Intermediate Flow Selection Process.

* Complete Flow

  A Complete Flow consists of all the packets that enter the Intermediate Flow Selection Process within the Flow timeout interval and that belong to the same Flow, per the definition of "Flow" in [RFC5470]. For this definition, only packets that arrive at the Intermediate Flow Selection Process are considered.

* Flow Position

  Flow Position is the position of a Flow Record within the Flow Cache.

* Flow Filtering

  Flow Filtering selects flows based on a deterministic function on the Flow Record content, Flow Selection State, external properties (e.g., ingress interface), or external events (e.g., violated Access Control List). If the relevant parts of the Flow Record content can already be observed at the packet level (e.g., Flow Keys from packet header fields), Flow Filtering can be performed at the packet level by Property Match Filtering, as described in [RFC5475].

* Hash-based Flow Filtering

  Hash-based Flow Filtering is a deterministic Flow filter function that selects flows based on a hash function. The hash function is calculated over parts of the Flow Record content or external properties that are called the Hash Domain. If the hash value falls into a predefined Hash Selection Range, the Flow is selected.

* Flow-state Dependent Intermediate Flow Selection Process

  The Flow-state dependent Intermediate Flow Selection Process is a selection function that selects or drops Flows based on the current Flow Selection State. The selection can be either deterministic, random, or non-uniform random.

* Flow-state Dependent Packet Selection

  Flow-state dependent packet selection is a selection function that selects or drops packets based on the current Flow Selection State. The selection can be either deterministic, random, or non-uniform random. Flow-state dependent packet selection can be used to implement a preference for the selection of packets belonging to specific Flows. For example, the selection probability of packets belonging to Flows that are already within the Flow Cache may be higher than for packets that have not been recorded yet.

* Flow Sampling

  Flow Sampling selects flows based on Flow Record sequence or arrival times (e.g., entry in Flow Cache, arrival time at Exporter or Mediator). The selection can be systematic (e.g., every n-th Flow) or based on a random function (e.g., select each Flow Record with probability p, or randomly select n out of N Flow Records).

3. Difference between Intermediate Flow Selection Process and Packet Selection

The Intermediate Flow Selection Process differs from packet selection as described in [RFC5475]. Packet selection techniques consider packets as the basic element, and the parent population consists of all packets observed at an Observation Point. In contrast to this, the basic elements in Flow selection are the Flows. The parent population consists of all observed Flows, and the Intermediate Flow Selection Process operates on the Flows. The major characteristics of the Intermediate Flow Selection Process are the following:

- The Intermediate Flow Selection Process takes Flows as basic elements. For packet selection, packets are considered as basic elements.

- The Intermediate Flow Selection Process typically takes place after Packet Classification, because the classification rules determine to which Flow a packet belongs. The Intermediate Flow Selection Process can be performed before Packet Classification. In that case, the Intermediate Flow Selection Process is based on the Flow Key (and also on a hash value over the Flow Key) but not

on characteristics that are only available after Packet Classification (e.g., Flow size, Flow duration). Packet selection can be applied before and after Packet Classification. As an example, packet selection before Packet Classification can be random packet selection, whereas packet selection after Packet Classification can be Flow-state dependent packet selection (as described in [RFC5475]).

- The Intermediate Flow Selection Process operates on Complete Flows. That means that after the Intermediate Flow Selection Process, either all packets of the Flow are kept or all packets of the Flow are discarded. That means that if the Intermediate Flow Selection Process is preceded by a packet selection process, the Complete Flow consists only of the packets that were not discarded during the packet selection.

There are some techniques that are difficult to unambiguously categorize into one of the categories. Here, some guidance is given on how to categorize such techniques:

- Techniques that can be considered as both packet selection and an Intermediate Flow Selection Process: some packet selection techniques result in the selection of Complete Flows and therefore can be considered as packet selection or as an Intermediate Flow Selection Process at the same time. An example is Property Match Filtering of all packets to a specific destination address. If Flows are defined based on destination addresses, such a packet selection also results in an Intermediate Flow Selection Process and can be considered as packet selection or as an Intermediate Flow Selection Process.

- Flow-state Dependent Packet Selection: there exist techniques that select packets based on the Flow state, e.g., based on the number of already observed packets belonging to the Flow. Examples of these techniques from the literature include "Sample and Hold" [EsVa01], "Fast Filtered Sampling" [MSZC10], and the "Sticky Sampling" algorithm presented in [MaMo02]. Such techniques can be used to influence which Flows are captured (e.g., increase the selection of packets belonging to large Flows) and reduce the number of Flows that need to be stored in the Flow Cache. Nevertheless, such techniques do not necessarily select Complete Flows, because they do not ensure that all packets of a selected Flow are captured. Therefore, Flow-state dependent packet selection techniques that do not ensure that either all or no packets of a Flow are selected, strictly speaking, have to be considered as packet selection techniques and not as Intermediate Flow Selection Process techniques.

4. Difference between Intermediate Flow Selection Process and Intermediate Selection Process

The Intermediate Flow Selection Process differs from the Intermediate Selection Process, since the Intermediate Flow Selection Process uses selection criteria that apply to a larger set of Flow information and properties than those used by the Intermediate Selection Process. The typical function of an Intermediate Selection Process is Property Match Filtering, which selects a Flow Record if the value of a specific field in the Flow Record matches a configured value or falls within a configured range. This means that the selection criteria used by an Intermediate Selection Process are evaluated only on Flow Record values. An Intermediate Flow Selection Process makes its decision on whether a Flow has to be selected or not by taking into account not only information related to the content of the Flow Record but also any Flow Selection State information or variable that can be used to select Flows in order to meet application requirements or resource constraints (e.g., Flow Cache occupancy, export link capacity). Examples include flow counters, Intermediate Flow Selection Process timeout intervals, and Flow Record time information.

5. Intermediate Flow Selection Process within the IPFIX Architecture

An Intermediate Flow Selection Process can be deployed at any of three places within the IPFIX architecture. As shown in Figure 1, the Intermediate Flow Selection Process can occur

1. in the Metering Process at the IPFIX Exporter

2. in the Exporting Process at the Collector

3. within a Mediator

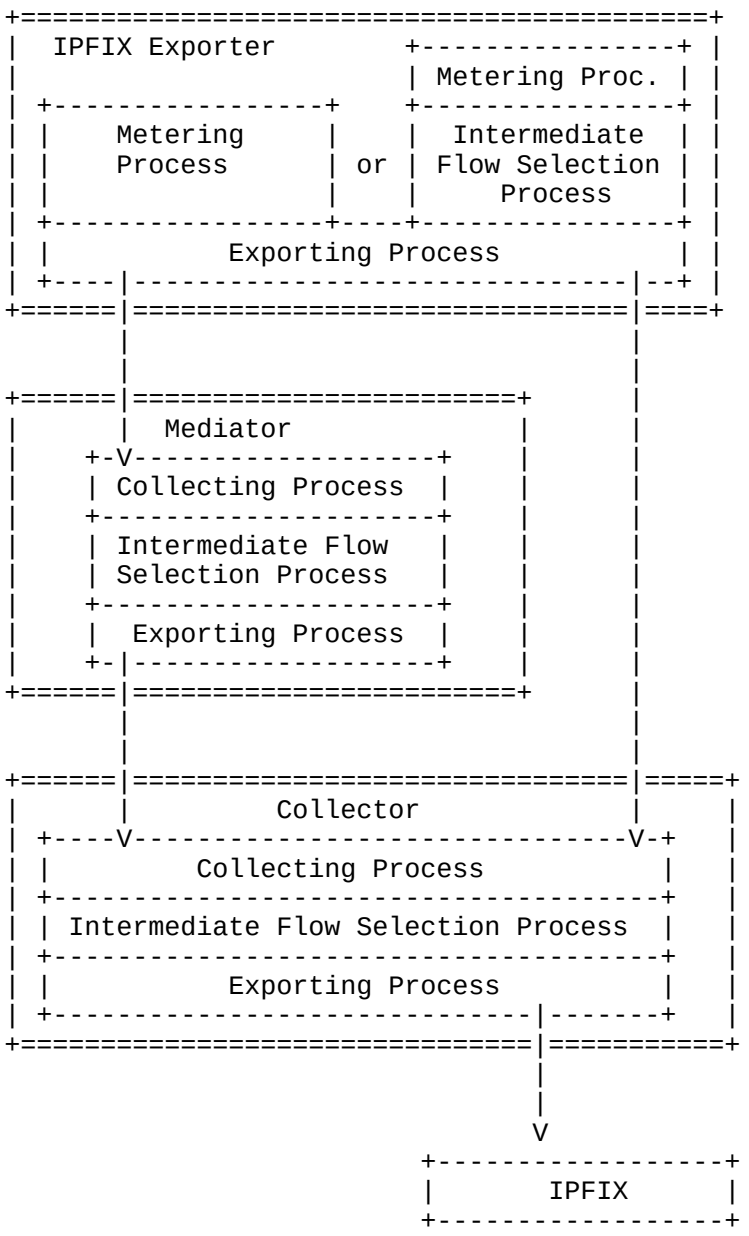

Figure 1: Potential Intermediate Flow Selection Process Locations

In contrast to packet selection, the Intermediate Flow Selection Process is always applied after the packets are classified into Flows.

5.1. Intermediate Flow Selection Process in the Metering Process

An Intermediate Flow Selection Process in the Metering Process uses packet information to update the Flow Records in the Flow Cache. The Intermediate Flow Selection Process, before Packet Classification, can be based on the Flow Key (and also on a hash value over the Flow Key) but not on characteristics that are only available after Packet Classification (e.g., Flow size, Flow duration). Here, an Intermediate Flow Selection Process is applied to reduce resources for all subsequent processes or to select specific Flows of interest in cases where such Flow characteristics are already observable at the packet level (e.g., Flows to specific IP addresses). In contrast, Flow-state dependent packet selection is a packet selection technique, because it does not necessarily select Complete Flows.

5.2. Intermediate Flow Selection Process in the Exporting Process

An Intermediate Flow Selection Process in the Exporting Process works on Flow Records and can therefore depend on Flow characteristics that are only visible after the classification of packets, such as Flow size and Flow duration. The Exporting Process may implement policies for exporting only a subset of the Flow Records that have been stored in the system's memory, in order to offload Flow export and Flow post-processing. An Intermediate Flow Selection Process in the Exporting Process may select only the subset of Flow Records that are of interest to the user's application or select only as many Flow Records as can be handled by the available resources (e.g., limited export link capacity).

5.3. Intermediate Flow Selection Process as a Function of the IPFIX Mediator

As shown in Figure 1, the Intermediate Flow Selection Process can be performed within an IPFIX Mediator [RFC6183]. The Intermediate Flow Selection Process takes a Flow Record stream as its input and selects Flow Records from a sequence based upon criteria-evaluated record values. The Intermediate Flow Selection Process can again apply an Intermediate Flow Selection Process technique to obtain Flows of interest to the application. Further, the Intermediate Flow Selection Process can base its selection decision on the correlation of data from different IPFIX Exporters, e.g., by only selecting Flows that were recorded on two or more IPFIX Exporters.

6. Intermediate Flow Selection Process Techniques

An Intermediate Flow Selection Process technique selects either all or none of the packets of a Flow; otherwise, the technique has to be considered as packet selection. A difference between Flow Filtering and Flow sampling is recognized.

6.1. Flow Filtering

Flow Filtering is a deterministic function on the IPFIX Flow Record content. If the relevant Flow characteristics are already observable at the packet level (e.g., Flow Keys), Flow Filtering can be applied before aggregation at the packet level. In order to be compliant with IPFIX, at least one of this document's Flow Filtering schemes MUST be implemented.

6.1.1. Property Match Filtering

Property Match Filtering is performed similarly to Property Match Filtering for packet selection as described in [RFC5475]. The difference is that Flow Record fields are used here, instead of packet fields, to derive the selection decision. Property Match Filtering is used to select a specific subset of the Flows that are of interest to a particular application (e.g., all Flows to a specific destination, all large Flows, etc.). Properties on which the filtering is based can be Flow Keys, Flow Timestamps, or Per-Flow Counters as described in [RFC7012]. Examples include the Flow size in bytes, the number of packets in the Flow, the observation time of the first or last packet, and the maximum packet length. An example of Property Match Filtering is to select Flows with more than a threshold number of observed octets. The selection criteria can be a specific value, a set of specific values, or an interval. For example, a Flow is selected if destinationIPv4Address and the total number of packets of the Flow equal two predefined values. An Intermediate Flow Selection Process using Property Match Filtering in the Metering Process relies on properties that are observable at the packet level (e.g., Flow Key). For example, a Flow is selected if sourceIPv4Address and sourceIPv4PrefixLength equal, respectively, two specific values.

An Intermediate Flow Selection Process using Property Match Filtering in the Exporting Process is based on properties that are only visible after Packet Classification, such as Flow size and Flow duration. An example is the selection of the largest Flows or a percentage of Flows with the longest lifetime. Another example is to select and remove from the Flow Cache the Flow Record with the lowest Flow volume per current Flow lifetime if the Flow Cache is full.

An Intermediate Flow Selection Process using Property Match Filtering within an IPFIX Mediator selects a Flow Record if the value of a specific field in the Flow Record equals a configured value or falls within a configured range [RFC6183].

6.1.2. Hash-Based Flow Filtering

Hash-based Flow Filtering uses a hash function h to map the Flow Key c onto a Hash Range R. A Flow is selected if the hash value h(c) is within the Hash Selection Range S, which is a subset of R. Hash-based Flow Filtering can be used to emulate a random sampling process but still enable the correlation between selected Flow subsets at different Observation Points. Hash-based Flow Filtering is similar to Hash-based packet selection and is in fact identical when Hash-based packet selection uses the Flow Key that defines the Flow as the hash input. Nevertheless, there may be the incentive to apply Hash-based Flow Filtering, but not at the packet level, in the Metering Process, for example, when the size of the selection range, and therefore the sampling probability, are dependent on the number of observed Flows. If Hash-based Flow Filtering is used to select the same subset of flows at different Observation Points, the Hash Domain MUST only include parts of the Flow Record content that are invariant on the Flow path. Refer also to the Trajectory Sampling application example of coordinated packet selection [RFC5475], which explains the hash-based filtering approach at the packet level.

6.2. Flow Sampling

Flow sampling operates on Flow Record sequence or arrival times. It can use either a systematic or a random function for the Intermediate Flow Selection Process. Flow sampling usually aims at the selection of a representative subset of all Flows in order to estimate characteristics of the whole set (e.g., mean Flow size in the network).

6.2.1. Systematic Sampling

Systematic sampling is a deterministic selection function. It may be a periodic selection of the N-th Flow Record that arrives at the Intermediate Flow Selection Process. Systematic sampling MAY be applied in the Metering Process. An example would be to create, besides the Flow Cache of selected Flows, an additional data structure that saves the Flow Key values of the Flows that are not selected. The selection of a Flow would then be based on the first packet of a Flow. Every time a packet belonging to a new Flow (which is not in the data structure of either the selected or non-selected Flows) arrives at the Observation Point, a counter is increased. If

the counter is increased to a multiple of N, a new Flow Cache entry is created; if the counter is not a multiple of N, the Flow Key value is added to the data structure for non-selected Flows.

Systematic sampling can also be time-based. Time-based systematic sampling is applied by only creating Flows that are observed between time-based start and stop triggers. The time interval may be applied at the packet level in the Metering Process or after aggregation at the Flow level, e.g., by selecting a Flow arriving at the Exporting Process every n seconds.

6.2.2. Random Sampling

Random Flow sampling is based on a random process that requires the calculation of random numbers. One can differentiate between n-out-of-N and probabilistic Flow sampling.

6.2.2.1. n-out-of-N Flow Sampling

In n-out-of-N Sampling, n elements are selected out of the parent population, which consists of N elements. One example would be to generate n different random numbers in the range [1,N] and select all Flows that have a Flow Position equal to one of the random numbers.

6.2.2.2. Probabilistic Flow Sampling

In probabilistic Sampling, the decision of whether or not a Flow is selected is made in accordance with a predefined selection probability. For probabilistic Sampling, the Sample Size can vary for different trials. The selection probability does not necessarily have to be the same for each Flow. Therefore, a difference between uniform probabilistic sampling (with the same selection probability for all Flows) and non-uniform probabilistic sampling (where the selection probability can vary for different Flows) is recognized. For non-uniform probabilistic Flow sampling, the sampling probability may be adjusted according to the Flow Record content. An example would be to increase the selection probability of large-volume Flows over small-volume Flows, as described in [DuLT01].

6.3. Flow-State Dependent Intermediate Flow Selection Process

The Flow-state dependent Intermediate Flow Selection Process can be a deterministic or random Intermediate Flow Selection Process, based on the Flow Record content and the Flow state that may be kept additionally for each of the Flows. External processes may update counters, bounds, and timers for each of the Flow Records, and the Intermediate Flow Selection Process utilizes this information for the selection decision. A review of Flow-state dependent Intermediate

Flow Selection Process techniques that aim at the selection of the most frequent items by keeping additional Flow state information can be found in [CoHa08]. The Flow-state dependent Intermediate Flow Selection Process can only be applied after packet aggregation, when a packet has been assigned to a Flow. The Intermediate Flow Selection Process then decides, based on the Flow state for each Flow, whether it is kept in the Flow Cache or not. Two Flow-state dependent Intermediate Flow Selection Process Algorithms are described here:

The Frequent algorithm [KaPS03] is a technique that aims at the selection of all flows that at least exceed a 1/k fraction of the Observed Packet Stream. The algorithm has only a Flow Cache of size k-1, and each Flow in the Flow Cache has an additional counter. The counter is incremented each time a packet belonging to the Flow in the Flow Cache is observed. If the observed packet does not belong to any Flow, all counters are decremented; if any of the Flow counters has a value of zero, the Flow is replaced with a Flow formed from the new packet.

Lossy counting is a selection technique that identifies all Flows whose packet count exceeds a certain percentage of the whole observed packet stream (e.g., 5% of all packets) with a certain estimation error e. Lossy counting separates the observed packet stream in windows of size N=1/e, where N is an amount of consecutive packets. For each observed Flow, an additional counter will be held in the Flow state. The counter is incremented each time a packet belonging to the Flow is observed, and all counters are decremented at the end of each window. Also, all Flows with a counter of zero are removed from the Flow Cache.

6.4. Flow-State Dependent Packet Selection

Flow-state dependent packet selection is not an Intermediate Flow Selection Process technique but a packet selection technique. Nevertheless, configuration and reporting parameters for this technique will be described in this document. An example is the "Sample and Hold" algorithm [EsVa01], which tries to implement a preference for large-volume Flows in the selection. When a packet arrives, it is selected when a Flow Record for this packet already exists. If there is no Flow Record, the packet is selected according to a certain probability that is dependent on the packet size.

7. Configuration of Intermediate Flow Selection Process Techniques

This section describes the configuration parameters of the Flow selection techniques presented above. It provides the basis for an information model to be adopted in order to configure the Intermediate Flow Selection Process within an IPFIX Device. The information model with the Information Elements (IEs) for Intermediate Flow Selection Process configuration is described together with the reporting IEs in Section 8. Table 1 gives an overview of the defined Intermediate Flow Selection Process techniques, where they can be applied, and what their input parameters are. Depending on where the Flow selection techniques are applied, different input parameters can be configured.

Location	Selection Technique	Selection Input
In the Metering Process	Flow-state Dependent Packet Selection	packet sampling probabilities, Flow Selection State, packet properties
In the Metering Process	Property Match Flow Filtering	Flow Record IEs, Selection Interval
In the Metering Process	Hash-based Flow Filtering	selection range, hash function, Flow Key, seed (optional)
In the Metering Process	Time-based Systematic Flow sampling	Flow Position (derived from arrival time of packets), Flow Selection State
In the Metering Process	Sequence-based Systematic Flow sampling	Flow Position (derived from packet position), Flow Selection State
In the Metering Process	Random Flow sampling	random number generator or list and packet position, Flow state
In the Exporting Process/ within the IPFIX Mediator	Property Match Flow Filtering	Flow Record content, filter function

In the Exporting Process/ within the IPFIX Mediator	Hash-based Flow Filtering	selection range, hash function, hash input (Flow Keys and other Flow properties)
In the Exporting Process/ within the IPFIX Mediator	Flow-state Dependent Intermediate Flow Selection Process	Flow state parameters, random number generator or list
In the Exporting Process/ within the IPFIX Mediator	Time-based Systematic Flow sampling	Flow arrival time, Flow state
In the Exporting Process/ within the IPFIX Mediator	Sequence-based Systematic Flow sampling	Flow Position, Flow state
In the Exporting Process/ within the IPFIX Mediator	Random Flow sampling	random number generator or list and Flow Position, Flow state

Table 1: Overview of Intermediate Flow Selection Process Techniques

7.1. Intermediate Flow Selection Process Parameters

This section defines what parameters are required to describe the most common Intermediate Flow Selection Process techniques.

Intermediate Flow Selection Process Parameters:

For Property Match Filtering:

- Information Element as specified in [IANA-IPFIX]):
  Specifies the Information Element that is used as the property in the filter expression. Section 8 specifies the Information Elements that MUST be exported by an Intermediate Flow Selection Process using Property Match Filtering.

- Selection Value or Value Interval:
  Specifies the value or interval of the filter expression. Packets and Flow Records that have a value equal to the Selection Value or within the Interval will be selected.

For Hash-based Flow Filtering:

- Hash Domain:
  Specifies the bits from the packet or Flow that are taken as the hash input to the hash function.

- Hash Function:
  Specifies the name of the hash function that is used to calculate the hash value. Possible hash functions are BOB [RFC5475], IP Shift-XOR (IPSX) [RFC5475], and CRC-32 [Bra75].

- Hash Selection Range:
  Flows that have a hash value within the Hash Selection Range are selected. The Hash Selection Range can be a value interval or arbitrary hash values within the Hash Range of the hash function.

- Random Seed or Initializer Value:
  Some hash functions require an initializing value. In order to make the selection decision more secure, one can choose a random seed that configures the hash function.

For Flow-state Dependent Intermediate Flow Selection Process:

- Frequency threshold:
  Specifies the frequency threshold s for Flow-state dependent Flow Selection techniques that try to find the most frequent items within a dataset. All Flows that exceed the defined threshold will be selected.

- Accuracy parameter:
  Specifies the accuracy parameter e for techniques that deal with the issue of mining frequent items in a dataset. The accuracy parameter defines the maximum error, i.e., no Flows that have a true frequency less than (s - e) N are selected, where s is the frequency threshold and N is the total number of packets.

The above list of parameters for Flow-state dependent Flow Selection techniques is suitable for the presented frequent item and lossy counting algorithms. Nevertheless, a variety of techniques exist with very specific parameters not defined here.

For Systematic time-based Flow sampling:

- Interval length (in usec):
  Defines the length of the sampling interval during which Flows are selected.

- Spacing (in usec):
  Defines the spacing in usec between the end of one sampling interval and the start of the next interval.

For Systematic count-based Flow sampling:

- Interval length:
  Defines the number of Flows that are selected within the sampling interval.

- Spacing:
  Defines the spacing, in number of observed Flows, between the end of one sampling interval and the start of the next interval.

For random n-out-of-N Flow sampling:

- Population Size N:
  The number of all Flows in the Population from which the sample is drawn.

- Sampling Size n:
  The number of Flows that are randomly drawn from the population N.

For probabilistic Flow sampling:

- Sampling probability p:
  Defines the probability by which each of the observed Flows is selected.

7.2. Description of Flow-State Dependent Packet Selection

The configuration of Flow-state dependent packet selection has not been described in [RFC5475]; therefore, the parameters are defined here:

For Flow-state Dependent Packet Selection:

- Packet selection probability per possible Flow state interval:
  Defines multiple {Flow interval, packet selection probability} value pairs that configure the sampling probability, depending on the current Flow state.

- Additional parameters:
  For the configuration of Flow-state dependent packet selection, additional parameters or packet properties may be required, e.g., the packet size [EsVa01].

8. Information Model for Intermediate Flow Selection Process Configuration and Reporting

This section specifies the Information Elements that MUST be exported by an Intermediate Flow Selection Process in order to support the interpretation of measurement results from Flow measurements. The information is mainly used to report how many packets and Flows have been observed in total and how many of them were selected. This helps, for instance, to calculate the Attained Selection Fraction (see also [RFC5476]), which is an important parameter for providing an accuracy statement. The IEs can provide reporting information about Flow Records, packets, or bytes. The reported metrics are the total number of elements and the number of selected elements. The number of dropped elements can be derived from this information.

Table 2 shows a list of Intermediate Flow Selection Process Information Elements:

ID	Name	ID	Name
301	selectionSequenceID	302	selectorID
390	flowSelectorAlgorithm	1	octetDeltaCount
391	flowSelectedOctetDeltaCount	2	packetDeltaCount
392	flowSelectedPacketDeltaCount	3	originalFlowsPresent
393	flowSelectedFlowDeltaCount	394	selectorIDTotalFlowsObserved
395	selectorIDTotalFlowsSelected	396	samplingFlowInterval
397	samplingFlowSpacing	309	samplingSize
310	samplingPopulation	311	samplingProbability
398	flowSamplingTimeInterval	399	flowSamplingTimeSpacing
326	digestHashValue	400	hashFlowDomain
329	hashOutputRangeMin	330	hashOutputRangeMax
331	hashSelectedRangeMin	332	hashSelectedRangeMax
333	hashDigestOutput	334	hashInitialiserValue
320	absoluteError	321	relativeError
336	upperCILimit	337	lowerCILimit
338	confidenceLevel		

Table 2: Intermediate Flow Selection Process Information Elements

9. IANA Considerations

9.1. Registration of Information Elements

IANA has registered the following IEs in the "IPFIX Information Elements" registry at http://www.iana.org/assignments/ipfix/.

9.1.1. flowSelectorAlgorithm

Description:

This Information Element identifies the Intermediate Flow Selection Process technique (e.g., Filtering, Sampling) that is applied by the Intermediate Flow Selection Process. Most of these techniques have parameters; configuration parameter(s) MUST be clearly specified. Further Information Elements are needed to fully specify packet selection with these methods and all their parameters. Further method identifiers may be added to the list below. It might be necessary to define new Information Elements to specify their parameters. The flowSelectorAlgorithm registry is maintained by IANA. New assignments for the registry will be administered by IANA, on a First Come First Served basis [RFC5226], subject to Expert Review [RFC5226]. Please note that the purpose of the flow selection techniques described in this document is the improvement of measurement functions as defined in the Introduction (Section 1). Before adding new flow selector algorithms, their intended purposes should be determined, especially if those purposes contradict any policies defined in [RFC2804]. The designated expert(s) should consult with the community if a request that runs counter to [RFC2804] is received. The registry can be updated when specifications of the new method(s) and any new Information Elements are provided. The group of experts must double-check the flowSelectorAlgorithm definitions and Information Elements with already-defined flowSelectorAlgorithm definitions and Information Elements for completeness, accuracy, and redundancy. Those experts will initially be drawn from the Working Group Chairs and document editors of the IPFIX and PSAMP Working Groups. The following identifiers for Intermediate Flow Selection Process Techniques are defined here:

ID	Technique	Parameters
1	Systematic count-based Sampling	flowSamplingInterval flowSamplingSpacing
2	Systematic time-based Sampling	flowSamplingTimeInterval flowSamplingTimeSpacing
3	Random n-out-of-N Sampling	samplingSize samplingPopulation
4	Uniform probabilistic Sampling	samplingProbability
5	Property Match Filtering	Information Element Value Range
	Hash-based Filtering	hashInitialiserValue hashFlowDomain
6	using BOB	hashSelectedRangeMin hashSelectedRangeMax
7	using IPSX	hashOutputRangeMin hashOutputRangeMax
8	using CRC	
9	Flow-state Dependent Intermediate Flow Selection Process	No agreed Parameters

Table 3: Intermediate Flow Selection Process Techniques

Abstract Data Type: unsigned16

ElementId: 390

Data Type Semantics: identifier

Status: current

9.1.2. flowSelectedOctetDeltaCount

Description:

This Information Element specifies the volume in octets of all Flows that are selected in the Intermediate Flow Selection Process since the previous report.

Abstract Data Type: unsigned64

ElementId: 391

Units: octets

Status: current

9.1.3. flowSelectedPacketDeltaCount

Description:

This Information Element specifies the volume in packets of all Flows that were selected in the Intermediate Flow Selection Process since the previous report.

Abstract Data Type: unsigned64

ElementId: 392

Units: packets

Status: current

9.1.4. flowSelectedFlowDeltaCount

Description:

This Information Element specifies the number of Flows that were selected in the Intermediate Flow Selection Process since the last report.

Abstract Data Type: unsigned64

ElementId: 393

Units: flows

Status: current

9.1.5. selectorIDTotalFlowsObserved

Description:

This Information Element specifies the total number of Flows observed by a Selector, for a specific value of SelectorID. This Information Element should be used in an Options Template scoped to the observation to which it refers. See Section 3.4.2.1 of the IPFIX protocol document [RFC7011].

Abstract Data Type: unsigned64

ElementId: 394

Units: flows

Status: current

9.1.6. selectorIDTotalFlowsSelected

Description:

This Information Element specifies the total number of Flows selected by a Selector, for a specific value of SelectorID. This Information Element should be used in an Options Template scoped to the observation to which it refers. See Section 3.4.2.1 of the IPFIX protocol document [RFC7011].

Abstract Data Type: unsigned64

ElementId: 395

Units: flows

Status: current

9.1.7. samplingFlowInterval

Description:

This Information Element specifies the number of Flows that are consecutively sampled. A value of 100 means that 100 consecutive Flows are sampled. For example, this Information Element may be used to describe the configuration of a systematic count-based Sampling Selector.

Abstract Data Type: unsigned64

ElementId: 396

Units: flows

Status: current

9.1.8. samplingFlowSpacing

Description:

This Information Element specifies the number of Flows between two "samplingFlowInterval"s. A value of 100 means that the next interval starts 100 Flows (which are not sampled) after the current "samplingFlowInterval" is over. For example, this Information Element may be used to describe the configuration of a systematic count-based Sampling Selector.

Abstract Data Type: unsigned64

ElementId: 397

Units: flows

Status: current

9.1.9. flowSamplingTimeInterval

Description:

This Information Element specifies the time interval in microseconds during which all arriving Flows are sampled. For example, this Information Element may be used to describe the configuration of a systematic time-based Sampling Selector.

Abstract Data Type: unsigned64

ElementId: 398

Units: microseconds

Status: current

9.1.10. flowSamplingTimeSpacing

Description:

This Information Element specifies the time interval in microseconds between two "flowSamplingTimeInterval"s. A value of 100 means that the next interval starts 100 microseconds (during which no Flows are sampled) after the current "flowsamplingTimeInterval" is over. For example, this Information Element may be used to describe the configuration of a systematic time-based Sampling Selector.

Abstract Data Type: unsigned64

ElementId: 399

Units: microseconds

Status: current

9.1.11. hashFlowDomain

Description:

This Information Element specifies the Information Elements that are used by the Hash-based Flow Selector as the Hash Domain.

Abstract Data Type: unsigned16

ElementId: 400

Data Type Semantics: identifier

Status: Current

9.2. Registration of Object Identifier

IANA has registered the following OID in the IPFIX-SELECTOR-MIB Functions subregistry at http://www.iana.org/assignments/smi-numbers according to the procedures set forth in [RFC6615].

Decimal	Name	Description	Reference
8	flowSelectorAlgorithm	This Object Identifier identifies the Intermediate Flow Selection Process technique (e.g., Filtering, Sampling) that is applied by the Intermediate Flow Selection Process	[RFC7014]

Table 4: Object Identifiers to Be Registered

10. Security and Privacy Considerations

Flow data exported by Exporting Processes, and collected by Collecting Processes, can be sensitive for privacy reasons and need to be protected. Privacy considerations for collected data are provided in [RFC7011].

Some of the described Intermediate Flow Selection Process techniques (e.g., Flow sampling, hash-based Flow Filtering) aim at the selection

of a representative subset of flows in order to estimate parameters of the population. An adversary may have incentives to influence the selection of flows, for example, to circumvent accounting or to avoid the detection of packets that are part of an attack.

Security considerations concerning the choice of a hash function for Hash-based packet selection have been discussed in Section 6.2.3 of [RFC5475] and are also appropriate for Hash-based Flow Selection. [RFC5475] discusses the possibility of crafting Packet Streams that are disproportionately selected or can be used to discover hash function parameters. It also describes vulnerabilities of different hash functions to these attacks and discusses practices to minimize these vulnerabilities.

For other sampling approaches, an adversary can gain knowledge about the start and stop triggers in time-based systematic Sampling, e.g., by sending test packets. This knowledge might allow adversaries to modify their send schedule in such a way that their packets are disproportionately selected or not selected. For random Sampling, an input to the encryption process, like the Initialization Vector of the CBC (Cipher Block Chaining) mode, should be used to prevent an adversary from predicting the selection decision [Dw01].

Further security threats can occur when Intermediate Flow Selection Process parameters are configured or communicated to other entities. The protocol(s) for the configuration and reporting of Intermediate Flow Selection Process parameters are out of scope for this document. Nevertheless, a set of initial requirements for future configuration and reporting protocols are stated below:

1. Protection against disclosure of configuration information: Intermediate Flow Selection Process configuration information describes the Intermediate Flow Selection Process and its parameters. This information can be useful to attackers. Attackers may craft packets that never fit the selection criteria in order to prevent Flows from being seen by the Intermediate Flow Selection Process. They can also craft a lot of packets that fit the selection criteria and overload or bias subsequent processes. Therefore, any transmission of configuration data (e.g., to configure a process or to report its actual status) should be protected by encryption.

2. Protection against modification of configuration information: Sending incorrect configuration information to the Intermediate Flow Selection Process can lead to a malfunction of the Intermediate Flow Selection Process. Additionally, reporting incorrect configuration information from the Intermediate Flow Selection Process to other processes can lead to incorrect

estimations at subsequent processes. Therefore, any protocol that transmits configuration information should prevent an attacker from modifying configuration information. Data integrity can be achieved by authenticating the data.

3. Protection against malicious nodes sending configuration information:
The remote configuration of Intermediate Flow Selection Process techniques should be protected against access by unauthorized nodes. This can be achieved by access control lists at the device that hosts the Intermediate Flow Selection Process (e.g., IPFIX Exporter, IPFIX Mediator, or IPFIX Collector) and by source authentication. The reporting of configuration data from an Intermediate Flow Selection Process has to be protected in the same way. That means that protocols that report configuration data from the Intermediate Flow Selection Process to other processes also need to protect against unauthorized nodes reporting configuration information.

The security threats that originate from communicating configuration information to and from Intermediate Flow Selection Processes cannot be assessed solely with the information given in this document. A further and more detailed assessment of security threats is necessary when a specific protocol for the configuration or reporting configuration data is proposed.

11. Acknowledgments

We would like to thank the IPFIX group, especially Brian Trammell, Paul Aitken, and Benoit Claise, for fruitful discussions and for proofreading the document.

12. References

12.1. Normative References

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

[RFC5475] Zseby, T., Molina, M., Duffield, N., Niccolini, S., and F. Raspall, "Sampling and Filtering Techniques for IP Packet Selection", RFC 5475, March 2009.

[RFC5476] Claise, B., Johnson, A., and J. Quittek, "Packet Sampling (PSAMP) Protocol Specifications", RFC 5476, March 2009.

[RFC6615] Dietz, T., Kobayashi, A., Claise, B., and G. Muenz, "Definitions of Managed Objects for IP Flow Information Export", RFC 6615, June 2012.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013.

[RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, September 2013.

12.2. Informative References

[Bra75] Brayer, K., "Evaluation of 32 Degree Polynomials in Error Detection on the SATIN IV Autovon Error Patterns", National Technical Information Service, August 1975.

[CoHa08] Cormode, G. and M. Hadjieleftheriou, "Finding Frequent Items in Data Streams", Proceedings of the 34th International Conference on Very Large DataBases (VLDB), Auckland, New Zealand, Volume 1, Issue 2, pages 1530-1541, August 2008.

[DuLT01] Duffield, N., Lund, C., and M. Thorup, "Charging from Sampled Network Usage", ACM SIGCOMM Internet Measurement Workshop (IMW) 2001, pages 245-256, San Francisco, CA, USA, November 2001.

[Dw01] Dworkin, M., "Recommendation for Block Cipher Modes of Operation - Methods and Techniques", NIST Special Publication 800-38A, December 2001.

[EsVa01] Estan, C. and G,. Varghese, "New Directions in Traffic Measurement and Accounting: Focusing on the Elephants, Ignoring the Mice", ACM SIGCOMM Internet Measurement Workshop (IMW) 2001, San Francisco, CA, USA, November 2001.

[IANA-IPFIX] IANA, "IP Flow Information Export (IPFIX) Entities Registry", <http://www.iana.org/assignments/ipfix/>.

[KaPS03] Karp, R., Papadimitriou, C., and S. Shenker, "A simple algorithm for finding frequent elements in sets and bags", ACM Transactions on Database Systems, Volume 28, pages 51-55, March 2003.

[MSZC10] Mai, J., Sridharan, A., Zang, H., and C. Chuah, "Fast Filtered Sampling", Computer Networks Volume 54, Issue 11, pages 1885-1898, ISSN 1389-1286, August 2010.

[MaMo02] Manku, G. and R. Motwani, "Approximate Frequency Counts over Data Streams", Proceedings of the 28th International Conference on Very Large DataBases (VLDB), Hong Kong, China, pages 346-357, August 2002.

[RFC2804] IAB and IESG, "IETF Policy on Wiretapping", RFC 2804, May 2000.

[RFC3917] Quittek, J., Zseby, T., Claise, B., and S. Zander, "Requirements for IP Flow Information Export (IPFIX)", RFC 3917, October 2004.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, March 2009.

[RFC6183] Kobayashi, A., Claise, B., Muenz, G., and K. Ishibashi, "IP Flow Information Export (IPFIX) Mediation: Framework", RFC 6183, April 2011.

Authors' Addresses

Salvatore D'Antonio
University of Napoli "Parthenope"
Centro Direzionale di Napoli Is. C4
Naples 80143
Italy

Phone: +39 081 5476766
EMail: salvatore.dantonio@uniparthenope.it

Tanja Zseby
CAIDA/FhG FOKUS
San Diego Supercomputer Center (SDSC)
University of California, San Diego (UCSD)
9500 Gilman Drive
La Jolla, CA 92093-0505
USA

EMail: tanja.zseby@tuwien.ac.at

Christian Henke
Tektronix Communications Berlin
Wohlrabedamm 32
Berlin 13629
Germany

Phone: +49 17 2323 8717
EMail: christian.henke@tektronix.com

Lorenzo Peluso
University of Napoli
Via Claudio 21
Napoli 80125
Italy

Phone: +39 081 7683821
EMail: lorenzo.peluso@unina.it