

Network Working Group
Request for Comments: 5475
Category: Standards Track

T. Zseby
Fraunhofer FOKUS
M. Molina
DANTE
N. Duffield
AT&T Labs - Research
S. Niccolini
NEC Europe Ltd.
F. Raspall
EPSC-UPC
March 2009

Sampling and Filtering Techniques for IP Packet Selection

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

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

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

Abstract

This document describes Sampling and Filtering techniques for IP packet selection. It provides a categorization of schemes and defines what parameters are needed to describe the most common selection schemes. Furthermore, it shows how techniques can be combined to build more elaborate packet Selectors. The document provides the basis for the definition of information models for configuring selection techniques in Metering Processes and for reporting the technique in use to a Collector.

Table of Contents

| | |
|---|----|
| 1. Introduction | 3 |
| 1.1. Conventions Used in This Document | 4 |
| 2. PSAMP Documents Overview | 4 |
| 3. Terminology | 4 |
| 3.1. Observation Points, Packet Streams, and Packet Content | 4 |
| 3.2. Selection Process | 5 |
| 3.3. Reporting | 7 |
| 3.4. Metering Process | 7 |
| 3.5. Exporting Process | 8 |
| 3.6. PSAMP Device | 8 |
| 3.7. Collector | 8 |
| 3.8. Selection Methods | 8 |
| 4. Categorization of Packet Selection Techniques | 11 |
| 5. Sampling | 12 |
| 5.1. Systematic Sampling | 13 |
| 5.2. Random Sampling | 14 |
| 5.2.1. n-out-of-N Sampling | 14 |
| 5.2.2. Probabilistic Sampling | 14 |
| 6. Filtering | 16 |
| 6.1. Property Match Filtering | 16 |
| 6.2. Hash-Based Filtering | 19 |
| 6.2.1. Application Examples for Coordinated Packet Selection | 19 |
| 6.2.2. Desired Properties of Hash Functions | 21 |
| 6.2.3. Security Considerations for Hash Functions | 22 |
| 6.2.4. Choice of Hash Function | 26 |
| 7. Parameters for the Description of Selection Techniques | 29 |
| 7.1. Description of Sampling Techniques | 30 |
| 7.2. Description of Filtering Techniques | 31 |
| 8. Composite Techniques | 34 |
| 8.1. Cascaded Filtering->Sampling or Sampling->Filtering | 34 |
| 8.2. Stratified Sampling | 34 |
| 9. Security Considerations | 35 |
| 10. Contributors | 36 |
| 11. Acknowledgments | 36 |

| | |
|--|----|
| 12. References | 36 |
| 12.1. Normative References | 36 |
| 12.2. Informative References | 36 |
| Appendix A. Hash Functions | 40 |
| A.1 IP Shift-XOR (IPSX) Hash Function..... | 40 |
| A.2 BOB Hash Function..... | 41 |

1. Introduction

There are two main drivers for the evolution in measurement infrastructures and their underlying technology. First, network data rates are increasing, with a concomitant growth in measurement data. Second, the growth is compounded by the demand of measurement-based applications for increasingly fine-grained traffic measurements. Devices that perform the measurements, require increasingly sophisticated and resource-intensive measurement capabilities, including the capture of packet headers or even parts of the payload, and classification for flow analysis. All these factors can lead to an overwhelming amount of measurement data, resulting in high demands on resources for measurement, storage, transfer, and post processing.

The sustained capture of network traffic at line rate can be performed by specialized measurement hardware. However, the cost of the hardware and the measurement infrastructure required to accommodate the measurements preclude this as a ubiquitous approach. Instead, some form of data reduction at the point of measurement is necessary.

This can be achieved by an intelligent packet selection through Sampling or Filtering. Another way to reduce the amount of data is to use aggregation techniques (not addressed in this document). The motivation for Sampling is to select a representative subset of packets that allow accurate estimates of properties of the unsampled traffic to be formed. The motivation for Filtering is to remove all packets that are not of interest. Aggregation combines data and allows compact pre-defined views of the traffic. Examples of applications that benefit from packet selection are given in [RFC5474]. Aggregation techniques are out of scope of this document.

1.1. 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].

2. PSAMP Documents Overview

This document is one out of a series of documents from the PSAMP group.

[RFC5474]: "A Framework for Packet Selection and Reporting" describes the PSAMP framework for 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.

RFC 5475 (this document): "Sampling and Filtering Techniques for IP Packet Selection" describes the set of packet selection techniques supported by PSAMP.

[RFC5476]: "Packet Sampling (PSAMP) Protocol Specifications" specifies the export of packet information from a PSAMP Exporting Process to a PSAMP Collecting Process.

[RFC5477]: "Information Model for Packet Sampling Exports" defines an information and data model for PSAMP.

3. Terminology

The PSAMP terminology defined here is fully consistent with all terms listed in [RFC5474] but includes additional terms required for the description of packet selection methods. An architecture overview and possible configurations of PSAMP elements can be found in [RFC5474]. PSAMP terminology also aims at consistency with terms used in [RFC3917]. The relationship between PSAMP and IPFIX terms is described in [RFC5474].

In the PSAMP documents, all defined PSAMP terms are written capitalized. This document uses the same convention.

3.1. Observation Points, Packet Streams, and Packet Content

* Observation Point

An Observation Point [RFC5101] is a location in the network where packets can be observed. Examples include:

- (i) A line to which a probe is attached;

- (ii) a shared medium, such as an Ethernet-based LAN;
- (iii) a single port of a router, or set of interfaces (physical or logical) of a router;
- (iv) an embedded measurement subsystem within an interface.

Note that one Observation Point may be a superset of several other Observation Points. For example, one Observation Point can be an entire line card. This would be the superset of the individual Observation Points at the line card's interfaces.

* Observed Packet Stream

The Observed Packet Stream is the set of all packets observed at the Observation Point.

* Packet Stream

A Packet Stream denotes a set of packets from the Observed Packet Stream that flows past some specified point within the Metering Process. An example of a Packet Stream is the output of the selection process. Note that packets selected from a stream, e.g., by Sampling, do not necessarily possess a property by which they can be distinguished from packets that have not been selected. For this reason, the term "stream" is favored over "flow", which is defined as a set of packets with common properties [RFC3917].

* Packet Content

The Packet Content denotes the union of the packet header (which includes link layer, network layer, and other encapsulation headers) and the packet payload. At some Observation Points, the link header information may not be available.

3.2. Selection Process

* Selection Process

A Selection Process takes the Observed Packet Stream as its input and selects a subset of that stream as its output.

* Selection State

A Selection Process may maintain state information for use by the Selection Process. At a given time, the Selection State may depend on packets observed at and before that time, and other variables. Examples include:

- (i) sequence numbers of packets at the input of Selectors;
- (ii) a timestamp of observation of the packet at the Observation Point;
- (iii) iterators for pseudorandom number generators;
- (iv) hash values calculated during selection;
- (v) indicators of whether the packet was selected by a given Selector.

Selection Processes may change portions of the Selection State as a result of processing a packet. Selection State for a packet is to reflect the state after processing the packet.

* Selector

A Selector defines what kind of action a Selection Process performs on a single packet of its input. If selected, the packet becomes an element of the output Packet Stream.

The Selector can make use of the following information in determining whether a packet is selected:

- (i) the Packet Content;
- (ii) information derived from the packet's treatment at the Observation Point;
- (iii) any Selection State that may be maintained by the Selection Process.

* Composite Selector

A Composite Selector is an ordered composition of Selectors, in which the output Packet Stream issuing from one Selector forms the input Packet Stream to the succeeding Selector.

- * Primitive Selector

A Selector is primitive if it is not a Composite Selector.

- * Selection Sequence

From all the packets observed at an Observation Point, only a few packets are selected by one or more Selectors. The Selection Sequence is a unique value per Observation Domain describing the Observation Point and the Selector IDs through which the packets are selected.

3.3. Reporting

- * Packet Reports

Packet Reports comprise a configurable subset of a packet's input to the Selection Process, including the Packet's Content, information relating to its treatment (for example, the output interface), and its associated Selection State (for example, a hash of the Packet's Content).

- * Report Interpretation

Report Interpretation comprises subsidiary information, relating to one or more packets, that is used for interpretation of their Packet Reports. Examples include configuration parameters of the Selection Process.

- * Report Stream

The Report Stream is the output of a Metering Process, comprising two distinguished types of information: Packet Reports and Report Interpretation.

3.4. Metering Process

A Metering Process selects packets from the Observed Packet Stream using a Selection Process, and produces as output a Report Stream concerning the selected packets.

The PSAMP Metering Process can be viewed as analogous to the IPFIX Metering Process [RFC5101], which produces Flow Records as its output, with the difference that the PSAMP Metering Process always contains a Selection Process. The relationship between PSAMP and IPFIX is further described in [RFC5477] and [RFC5474].

3.5. Exporting Process

* Exporting Process

An Exporting Process sends, in the form of Export Packets, the output of one or more Metering Processes to one or more Collectors.

* Export Packet

An Export Packet is a combination of Report Interpretations and/or one or more Packet Reports that are bundled by the Exporting Process into an Export Packet for exporting to a Collector.

3.6. PSAMP Device

* PSAMP Device

A PSAMP Device is a device hosting at least an Observation Point, a Metering Process (which includes a Selection Process), and an Exporting Process. Typically, corresponding Observation Point(s), Metering Process(es), and Exporting Process(es) are colocated at this device, for example, at a router.

3.7. Collector

* Collector

A Collector receives a Report Stream exported by one or more Exporting Processes. In some cases, the host of the Metering and/or Exporting Processes may also serve as the Collector.

3.8. Selection Methods

* Filtering

A filter is a Selector that selects a packet deterministically based on the Packet Content, or its treatment, or functions of these occurring in the Selection State. Two examples are:

- (i) Property Match Filtering: A packet is selected if a specific field in the packet equals a predefined value.
- (ii) Hash-based Selection: A Hash Function is applied to the Packet Content, and the packet is selected if the result falls in a specified range.

* Sampling

A Selector that is not a filter is called a Sampling operation. This reflects the intuitive notion that if the selection of a packet cannot be determined from its content alone, there must be some type of Sampling taking place. Sampling operations can be divided into two subtypes:

- (i) Content-independent Sampling, which does not use Packet Content in reaching Sampling decisions. Examples include systematic Sampling, and uniform pseudorandom Sampling driven by a pseudorandom number whose generation is independent of Packet Content. Note that in content-independent Sampling, it is not necessary to access the Packet Content in order to make the selection decision.
- (ii) Content-dependent Sampling, in which the Packet Content is used in reaching selection decisions. An application is pseudorandom selection according to a probability that depends on the contents of a packet field, e.g., Sampling packets with a probability dependent on their TCP/UDP port numbers. Note that this is not a Filter.

* Hash Domain

A Hash Domain is a subset of the Packet Content and the packet treatment, viewed as an N-bit string for some positive integer N.

* Hash Range

A Hash Range is a set of M-bit strings for some positive integer M that defines the range of values that the result of the hash operation can take.

* Hash Function

A Hash Function defines a deterministic mapping from the Hash Domain into the Hash Range.

* Hash Selection Range

A Hash Selection Range is a subset of the Hash Range. The packet is selected if the action of the Hash Function on the Hash Domain for the packet yields a result in the Hash Selection Range.

- * Hash-based Selection

A Hash-based Selection is Filtering specified by a Hash Domain, a Hash Function, a Hash Range, and a Hash Selection Range.

- * Approximative Selection

Selectors in any of the above categories may be approximated by operations in the same or another category for the purposes of implementation. For example, uniform pseudorandom Sampling may be approximated by Hash-based Selection, using a suitable Hash Function and Hash Domain. In this case, the closeness of the approximation depends on the choice of Hash Function and Hash Domain.

- * Population

A Population is a Packet Stream or a subset of a Packet Stream. A Population can be considered as a base set from which packets are selected. An example is all packets in the Observed Packet Stream that are observed within some specified time interval.

- * Population Size

The Population Size is the number of all packets in the Population.

- * Sample Size

The Sample Size is a number of packets selected from the Population by a Selector.

- * Configured Selection Fraction

The Configured Selection Fraction is the expected ratio of the Sample Size to the Population Size, as based on the configured selection parameters.

- * Attained Selection Fraction

The Attained Selection Fraction is the ratio of the actual Sample Size to the Population Size. For some Sampling methods, the Attained Selection Fraction can differ from the Configured Selection Fraction due to, for example, the inherent statistical variability in Sampling decisions of probabilistic Sampling and Hash-based Selection. Nevertheless, for large Population Sizes and properly configured Selectors, the Attained Selection Fraction usually approaches the Configured Selection Fraction.

4. Categorization of Packet Selection Techniques

Packet selection techniques generate a subset of packets from an Observed Packet Stream at an Observation Point. We distinguish between Sampling and Filtering.

Sampling is targeted at the selection of a representative subset of packets. The subset is used to infer knowledge about the whole set of observed packets without processing them all. The selection can depend on packet position, and/or on Packet Content, and/or on (pseudo) random decisions.

Filtering selects a subset with common properties. This is used if only a subset of packets is of interest. The properties can be directly derived from the Packet Content, or depend on the treatment given by the router to the packet. Filtering is a deterministic operation. It depends on Packet Content or router treatment. It never depends on packet position or on (pseudo) random decisions.

Note that a common technique to select packets is to compute a Hash Function on some bits of the packet header and/or content and to select it if the hash value falls in the Hash Selection Range. Since hashing is a deterministic operation on the Packet Content, it is a Filtering technique according to our categorization. Nevertheless, Hash Functions are sometimes used to emulate random Sampling. Depending on the chosen input bits, the Hash Function, and the Hash Selection Range, this technique can be used to emulate the random selection of packets with a given probability p . It is also a powerful technique to consistently select the same packet subset at multiple Observation Points [DuGr00].

The following table gives an overview of the schemes described in this document and their categorization. X means that the characteristic applies to the selection scheme. (X) denotes schemes for which content-dependent and content-independent variants exist. For instance, Property Match Filtering is typically based on Packet Content and therefore is content dependent. But as explained in Section 6.1, it may also depend on router state and then would be independent of the content. It easily can be seen that only schemes with both properties, content dependence and deterministic selection, are considered as Filters.

| Selection Scheme | Deterministic Selection | Content - Dependent | Category |
|-------------------------------|-------------------------|---------------------|-----------|
| Systematic Count-based | X | - | Sampling |
| Systematic Time-based | X | - | Sampling |
| Random n-out-of-N | - | - | Sampling |
| Random uniform probabilistic | - | - | Sampling |
| Random non-uniform probabil. | - | (X) | Sampling |
| Random non-uniform Flow-State | - | (X) | Sampling |
| Property Match Filtering | X | (X) | Filtering |
| Hash function | X | X | Filtering |

The categorization just introduced is mainly useful for the definition of an information model describing Primitive Selectors. More complex selection techniques can be described through the composition of cascaded Sampling and Filtering operations. For example, a packet selection that weights the selection probability on the basis of the packet length can be described as a cascade of a Filtering and a Sampling scheme. However, this descriptive approach is not intended to be rigid: if a common and consolidated selection practice turns out to be too complex to be described as a composition of the mentioned building blocks, an ad hoc description can be specified instead and added as a new scheme to the information model.

5. Sampling

The deployment of Sampling techniques aims at the provisioning of information about a specific characteristic of the parent Population at a lower cost than a full census would demand. In order to plan a suitable Sampling strategy, it is therefore crucial to determine the needed type of information and the desired degree of accuracy in advance.

First of all, it is important to know the type of metric that should be estimated. The metric of interest can range from simple packet counts [JePP92] up to the estimation of whole distributions of flow characteristics (e.g., packet sizes) [ClPB93].

Second, the required accuracy of the information and with this, the confidence that is aimed at, should be known in advance. For instance, for usage-based accounting the required confidence for the estimation of packet counters can depend on the monetary value that corresponds to the transfer of one packet. That means that a higher confidence could be required for expensive packet flows (e.g., premium IP service) than for cheaper flows (e.g., best effort). The accuracy requirements for validating a previously agreed quality can also vary extremely with the customer demands. These requirements are usually determined by the service level agreement (SLA).

The Sampling method and the parameters in use must be clearly communicated to all applications that use the measurement data. Only with this knowledge a correct interpretation of the measurement results can be ensured.

Sampling methods can be characterized by the Sampling algorithm, the trigger type used for starting a Sampling interval, and the length of the Sampling interval. These parameters are described here in detail. The Sampling algorithm describes the basic process for selection of samples. In accordance to [AmCa89] and [ClPB93], we define the following basic Sampling processes.

5.1. Systematic Sampling

Systematic Sampling describes the process of selecting the start points and the duration of the selection intervals according to a deterministic function. This can be for instance the periodic selection of every k-th element of a trace but also the selection of all packets that arrive at predefined points in time. Even if the selection process does not follow a periodic function (e.g., if the time between the Sampling intervals varies over time), we consider this as systematic Sampling as long as the selection is deterministic.

The use of systematic Sampling always involves the risk of biasing the results. If the systematics in the Sampling process resemble systematics in the observed stochastic process (occurrence of the characteristic of interest in the network), there is a high probability that the estimation will be biased. Systematics in the observed process might not be known in advance.

Here only equally spaced schemes are considered, where triggers for Sampling are periodic, either in time or in packet count. All packets occurring in a selection interval (either in time or packet count) beyond the trigger are selected.

Systematic count-based

In systematic count-based Sampling, the start and stop triggers for the Sampling interval are defined in accordance to the spatial packet position (packet count).

Systematic time-based

In systematic time-based Sampling, time-based start and stop triggers are used to define the Sampling intervals. All packets are selected that arrive at the Observation Point within the time intervals defined by the start and stop triggers (i.e., arrival time of the packet is larger than the start time and smaller than the stop time).

Both schemes are content-independent selection schemes. Content-dependent deterministic Selectors are categorized as filters.

5.2. Random Sampling

Random Sampling selects the starting points of the Sampling intervals in accordance to a random process. The selection of elements is an independent experiment. With this, unbiased estimations can be achieved. In contrast to systematic Sampling, random Sampling requires the generation of random numbers. One can differentiate two methods of random Sampling: n-out-of-N Sampling and probabilistic Sampling.

5.2.1. n-out-of-N Sampling

In n-out-of-N Sampling, n elements are selected out of the parent Population that consists of N elements. One example would be to generate n different random numbers in the range [1,N] and select all packets that have a packet position equal to one of the random numbers. For this kind of Sampling, the Sample Size n is fixed.

5.2.2. Probabilistic Sampling

In probabilistic Sampling, the decision whether or not an element is selected is made in accordance to a predefined selection probability. An example would be to flip a coin for each packet and select all packets for which the coin showed the head. For this kind of Sampling, the Sample Size can vary for different trials. The selection probability does not necessarily have to be the same for each packet. Therefore, we distinguish between uniform probabilistic Sampling (with the same selection probability for all packets) and

non-uniform probabilistic Sampling (where the selection probability can vary for different packets).

5.2.2.1. Uniform Probabilistic Sampling

For Uniform Probabilistic Sampling, packets are selected independently with a uniform probability p . This Sampling can be count-driven, and is sometimes referred to as geometric random Sampling, since the difference in count between successive selected packets is an independent random variable with a geometric distribution of mean $1/p$. A time-driven analog, exponential random Sampling, has the time between triggers exponentially distributed.

Both geometric and exponential random Sampling are examples of what is known as additive random Sampling, defined as Sampling where the intervals or counts between successive samples are independent identically distributed random variables.

5.2.2.2. Non-Uniform Probabilistic Sampling

This is a variant of Probabilistic Sampling in which the Sampling probabilities can depend on the selection process input. This can be used to weight Sampling probabilities in order, e.g., to boost the chance of Sampling packets that are rare but are deemed important. Unbiased estimators for quantitative statistics are recovered by re-normalization of sample values; see [HT52].

5.2.2.3. Non-Uniform Flow State Dependent Sampling

Another type of Sampling that can be classified as probabilistic Non-Uniform is closely related to the flow concept as defined in [RFC3917], and it is only used jointly with a flow monitoring function (IPFIX Metering Process). Packets are selected, dependent on a Selection State. The point, here, is that the Selection State is determined also by the state of the flow the packet belongs to and/or by the state of the other flows currently being monitored by the associated flow monitoring function. An example for such an algorithm is the "sample and hold" method described in [EsVa01]:

- If a packet accounts for a Flow Record that already exists in the IPFIX flow recording process, it is selected (i.e., the Flow Record is updated).
- If a packet doesn't account for any existing Flow Record, it is selected with probability p . If it has been selected, a new Flow Record has to be created.

A further algorithm that fits into the category of non-uniform flow state dependent Sampling is described in [Moli03].

This type of Sampling is content dependent because the identification of the flow the packet belongs to requires analyzing part of the Packet Content. If the packet is selected, then it is passed as an input to the IPFIX monitoring function (this is called "Local Export" in [RFC5474]). Selecting the packet depending on the state of a flow cache is useful when memory resources of the flow monitoring function are scarce (i.e., there is no room to keep all the flows that have been scheduled for monitoring).

5.2.2.4. Configuration of Non-Uniform Probabilistic and Flow State Sampling

Many different specific methods can be grouped under the terms non-uniform probabilistic and flow state Sampling. Dependent on the Sampling goal and the implemented scheme, a different number and type of input parameters are required to configure such a scheme.

Some concrete proposals for such methods exist from the research community (e.g., [EsVa01], [DuLT01], [Moli03]). Some of these proposals are still in an early stage and need further investigations to prove their usefulness and applicability. It is not our aim to indicate preference among these methods. Instead, we only describe here the basic methods and leave the specification of explicit schemes and their parameters up to vendors (e.g., as an extension of the information model).

6. Filtering

Filtering is the deterministic selection of packets based on the Packet Content, the treatment of the packet at the Observation Point, or deterministic functions of these occurring in the Selection State. The packet is selected if these quantities fall into a specified range. The role of Filtering, as the word itself suggest, is to separate all the packets having a certain property from those not having it. A distinguishing characteristic from Sampling is that the selection decision does not depend on the packet position in time or in space, or on a random process.

We identify and describe in the following two Filtering techniques.

6.1. Property Match Filtering

With this Filtering method, a packet is selected if specific fields within the packet and/or properties of the router state equal a predefined value. Possible filter fields are all IPFIX flow

attributes specified in [RFC5102]. Further fields can be defined by proposing new information elements or defining vendor-specific extensions.

A packet is selected if Field=Value. Masks and ranges are only supported to the extent to which [RFC5102] allows them, e.g., by providing explicit fields like the netmasks for source and destination addresses.

AND operations are possible by concatenating filters, thus producing a composite selection operation. In this case, the ordering in which the Filtering happens is implicitly defined (outer filters come after inner filters). However, as long as the concatenation is on filters only, the result of the cascaded filter is independent from the order, but the order may be important for implementation purposes, as the first filter will have to work at a higher rate. In any case, an implementation is not constrained to respect the filter ordering, as long as the result is the same, and it may even implement the composite Filtering in one single step.

OR operations are not supported with this basic model. More sophisticated filters (e.g., supporting bitmasks, ranges, or OR operations) can be realized as vendor-specific schemes.

All IPFIX flow attributes defined in [RFC5102] can be used for Property Match Filtering. Further information elements can be easily defined. Property match operations should be available for different protocol portions of the packet header:

- (i) IP header (excluding options in IPv4, stacked headers in IPv6)
- (ii) transport protocol header (e.g., TCP, UDP)
- (iii) encapsulation headers (e.g., the MPLS label stack, if present)

When the PSAMP Device offers Property Match Filtering, and, in its usual capacity other than in performing PSAMP functions, identifies or processes information from IP, transport protocol or encapsulation protocols, then the information should be made available for Filtering. For example, when a PSAMP Device routes based on destination IP address, that field should be made available for Filtering. Conversely, a PSAMP Device that does not route is not expected to be able to locate an IP address within a packet, or make it available for Filtering, although it may do so.

Since packet encryption conceals the real values of encrypted fields, Property Match Filtering must be configurable to ignore encrypted packets, when detected.

The Selection Process may support Filtering based on the properties of the router state:

- (i) Ingress interface at which a packet arrives equals a specified value
- (ii) Egress interface to which a packet is routed to equals a specified value
- (iii) Packet violated Access Control List (ACL) on the router
- (iv) Failed Reverse Path Forwarding (RPF)
- (v) Failed Resource Reservation Protocol (RSVP)
- (vi) No route found for the packet
- (vii) Origin Border Gateway Protocol (BGP) Autonomous System (AS) [RFC4271] equals a specified value or lies within a given range
- (viii) Destination BGP AS equals a specified value or lies within a given range

Packets that match the failed Reverse Path Forwarding (RPF) condition are packets for which ingress Filtering failed as defined in [RFC3704].

Packets that match the failed Resource Reservation Protocol (RSVP) condition are packets that do not fulfill the RSVP specification as defined in [RFC2205].

Router architectural considerations may preclude some information concerning the packet treatment being available at line rate for selection of packets. For example, the Selection Process may not be implemented in the fast path that is able to access router state at line rate. However, when Filtering follows Sampling (or some other selection operation) in a Composite Selector, the rate of the Packet Stream output from the sampler and input to the filter may be sufficiently slow that the filter could select based on router state.

6.2. Hash-Based Filtering

A Hash Function h maps the Packet Content c , or some portion of it, onto a Hash Range R . The packet is selected if $h(c)$ is an element of S , which is a subset of R called the Hash Selection Range. Thus, Hash-Based selection is a particular case of Filtering. The object is selected if c is in $\text{inv}(h(S))$. But for desirable Hash Functions, the inverse image $\text{inv}(h(S))$ will be extremely complex, and hence h would not be expressible as, say, a Property Match Filter or a simple combination of these.

Hash-based Selection is mainly used to realize a coordinated packet selection. That means that the same packets are selected at different Observation Points. This is useful for instance to observe the path (trajectory) that a packet took through the network or to apply packet selection to passive one-way measurements.

A prerequisite for the method to work and to ensure interoperability is that the same Hash Function with the same parameters (e.g., input vector) is used at the Observation Points.

A consistent packet selection is also possible with Property Match Filtering. Nevertheless, Hash-based Selection can be used to approximate a random selection. The desired statistical properties are discussed in Section 6.2.2.

In the following subsections, we give some application examples for coordinated packet selection.

6.2.1. Application Examples for Coordinated Packet Selection

6.2.1.1. Trajectory Sampling

Trajectory Sampling is the consistent selection of a subset of packets at either all of a set of Observation Points or none of them. Trajectory Sampling is realized by Hash-based Selection if all Observation Points in the set use a common Hash Function, Hash Domain, and Selection Range. The Hash Domain comprises all or part of the Packet Content that is invariant along the packet path. Fields such as Time-to-Live, which is decremented per hop, and header CRC [RFC1624], which is recalculated per hop, are thus excluded from the Hash Domain. The Hash Domain needs to be wider than just a flow key, if packets are to be selected quasi-randomly within flows.

The trajectory (or path) followed by a packet is reconstructed from PSAMP reports on it that reach a Collector. Reports on a given packet originating from different Observation Points are associated by matching a label from the reports. The label may comprise that

portion of the invariant Packet Content that is reported, or possibly some digest of the invariant Packet Content that is inserted into the packet report at the Observation Point. Such a digest may be constructed by applying a second Hash Function to the invariant Packet Content. The reconstruction of trajectories and methods for dealing with possible ambiguities due to label collisions (identical labels reported for different packets) and potential loss of reports in transmission are dealt with in [DuGr00], [DuGG02], and [DuGr04].

Applications of trajectory Sampling include (i) estimation of the network path matrix, i.e., the traffic intensities according to network path, broken down by flow key; (ii) detection of routing loops, as indicated by self-intersecting trajectories; (iii) passive performance measurement: prematurely terminating trajectories indicate packet loss, packet one-way delay can be determined if reports include (synchronized) timestamps of packet arrival at the Observation Point; and (iv) network attack tracing, of the actual paths taken by attack packets with spoofed source addresses.

6.2.1.2. Passive One-Way Measurements

Coordinated packet selection can be applied for instance to one-way delay measurements in order to reduce the required resources. In one-way delay measurements, packets are collected at different Observation Points in the network. A packet digest is generated for each packet that helps to identify the packet. The packet digest and the arrival time of the packet at the Observation Point are reported to a process that calculates the delay. The delay is calculated by subtracting the arrival time of the same packet at the Observation Points (e.g., [ZsZC01]). With high data rates, capturing all packets can require a lot of resources for storage, transfer, and processing. To reduce resource consumption, packet selection methods can be applied. But for such selection techniques, it has to be ensured that the same packets are collected at different Observation Points. Hash-based Selection provides this feature.

6.2.1.3. Generation of Pseudorandom Numbers

Although pseudorandom number generators with well-understood properties have been developed, they may not be the method of choice in settings where computational resources are scarce. A convenient alternative is to use Hash Functions of Packet Content as a source of randomness. The hash (suitably re-normalized) is a pseudorandom variate in the interval $[0,1]$. Other schemes may use packet fields in iterators for pseudorandom numbers. However, the statistical properties of an ideal packet selection law (such as independent

Sampling for different packets, or independence on Packet Content) may not be exactly rendered by an implementation, but only approximately so.

Use of Packet Content to generate pseudorandom variates shares with non-uniform probabilistic Sampling (see Section 5.2.2.2 above) the property that selection decisions depend on Packet Content. However, there is a fundamental difference between the two. In the former case, the content determines pseudorandom variates. In the latter case, the content only determines the selection probabilities: selection could then proceed, e.g., by use of random variates obtained by an independent pseudorandom number generator.

6.2.2. Desired Properties of Hash Functions

Here we formulate desired properties for Hash Functions. For this, we have to distinguish whether a Hash Function is used for packet selection or just as a packet digest. The main focus of this document is on packet selection. Nevertheless, we also provide some requirements for the use of Hash Functions as packet digest.

First of all, we need to define suitable input fields from the packet. In accordance to [DuGr00], input field should be:

- invariant on the path
- variable among packets

Only if the input fields are the same at different Observation Points is it possible to recognize the packet. The input fields should be variable among packets in order to distribute the hash results over the selection range.

6.2.2.1. Requirements for Packet Selection

In accordance to considerations in [MoND05] and [Henk08], we define the following desired properties of Hash Functions used for packet selection:

- (i) Speed: The Hash Function has to be applied to each packet that traverses the Observation Point. Therefore, it has to be fast in order to cope with the high packet rates. In the ideal case, the hash operation should not influence the performance on the PSAMP Device.

- (ii) Uniformity: The Hash Function h should have good mixing properties, in the sense that small changes in the input (e.g., the flipping of a single bit) cause large changes in the output (many bits change). Then any local clump of values of c is spread widely over R by h , and so the distribution of $h(c)$ is fairly uniform even if the distribution of c is not. Then the Attained Selection Fraction gets close to the Configured Selection Fraction ($\#S/\#R$), which can be tuned by choice of S .
- (iii) Unbiasedness: The selection decision should be as independent of packet attributes as possible. The set of selected packets should not be biased towards a specific type of packets.
- (iv) Representativeness of sample: The sample should be as representative as possible for the observed traffic.
- (v) Non-linearity: The function should not be linear. This increases the mixing properties (uniformity criterion). In addition to this, it decreases the predictability of the output and therefore the vulnerabilities against attacks.
- (vi) Robustness against vulnerabilities: The Hash Function should be robust against attacks. Potential vulnerabilities are described in Section 6.2.3.

6.2.2.2. Requirements for Packet Digesting

For digesting Packet Content for inclusion in a reported label, the most important property is a low collision frequency. A secondary requirement is the ability to accept variable-length input, in order to allow inclusion of maximal amount of packet as input. Execution speed is of secondary importance, since the digest need only be formed from selected packets.

6.2.3. Security Considerations for Hash Functions

A concern for Hash-based Selection is whether some large set of related packets could be disproportionately sampled, i.e., that the Attained Selection Fraction is significantly different from the Configured Selection Fraction. This can happen either

- (i) through unanticipated behavior in the Hash Function, or
- (ii) because the packets had been deliberately crafted to have this property.

The first point underlines the importance of using a Hash Function with good mixing properties. For this, the statistical properties of candidate Hash Functions need to be evaluated. Since the hash output depends on the traffic mix, the evaluation should be done preferably on up-to-date packet traces from the network in which the Hash-based Selection will be deployed.

However, Hash Functions that perform well on typical traffic may not be sufficiently strong to withstand attacks specifically targeted against them. Such potential attacks have been described in [GoRe07].

In the following subsections, we point out different potential attack scenarios. We encourage the use of standardized Hash Functions. Therefore, we assume that the Hash Function itself is public and hence known to an attacker.

Nevertheless, we also assume the possibility of using a private input parameter for the Hash Function that is kept secret. Such an input parameter can for instance be attached to the hash input before the hash operation is applied. With this, at least parts of the hash operation remain secret.

For the attack scenarios, we assume that an attacker uses its knowledge of the Hash Function to craft packets that are then dispatched, either as the attack itself or to elicit further information that can be used to refine the attack.

Two scenarios are considered. In the first scenario, the attacker has no knowledge about whether or not the crafted packets are selected. In the second scenario, the attacker uses some knowledge of Sampling outcomes. The means by which this might be acquired is discussed below. Some additional attacks that involve tampering with Export Packets in transit, as opposed to attacking the PSAMP Device, are discussed in [GoRe07].

6.2.3.1. Vulnerabilities of Hash-Based Selection without Knowledge of Selection Outcomes

- (i) The Hash Function does not use a private parameter.

If no private input parameter is used, potential attackers can easily calculate which packets result in which hash values. If the selection range is public, an attacker can craft packets whose selection properties are known in advance. If the selection range is private, an attacker cannot determine whether a crafted packet is selected. However, by computing the hash on different trial crafted packets, and selecting those yielding a given hash value, the

attacker can construct an arbitrarily large set of distinct packets with a common selection properties, i.e., packets that will be either all selected or all not selected. This can be done whatever the strength of the Hash Function.

(ii) The Hash Function is not cryptographically strong.

If the Hash Function is not cryptographically strong, it may be possible to construct sequences of distinct packets with the common selection property even if a private parameter is used.

An example is the standard CRC-32 Hash Function used with a private modulus (but without a private string post-pended to the input). It has weak mixing properties for low-order bits. Consequently, simply by incrementing the hash input, one obtains distinct packets whose hashes mostly fall in a narrow range, and hence are likely commonly selected; see [GoRe07].

Suitable parameterization of the Hash Function can make such attacks more difficult. For example, post-pending a private string to the input before hashing with CRC-32 will give stronger mixing properties over all bits of the input. However, with a Hash Function, such as CRC-32, that is not cryptographically strong, the possibility of discovering a method to construct packet sets with the common selected property cannot be ruled out, even when a private modulus or post-pended string is used.

6.2.3.2. Vulnerabilities of Hash-Based Selection Using Knowledge of Selection Outcomes

Knowledge of the selection outcomes of crafted packets can be used by an attacker to more easily construct sets of packets that are disproportionately sampled and/or are commonly selected. For this, the attacker does not need any a priori knowledge about the Hash Function or selection range.

There are several ways an attacker might acquire this knowledge about the selection outcome:

- (i) Billing Reports: If samples are used for billing purposes, then the selection outcomes of packets may be able to be inferred by correlating a crafted Packet Stream with the billing reports that it generates. However, the rate at which knowledge of selection outcomes can be acquired depends on the temporal and spatial granularity of the billing reports; being slower the more aggregated the reports are.

- (ii) Feedback from an Intrusion Detection System: e.g., a botmaster adversary learns if his packets were detected by the intrusion detection system by seeing if one of his bots is blocked by the network.
- (iii) Observation of the Report Stream: Export Packets sent across a public network may be eavesdropped on by an adversary. Encryption of the Export Packets provides only a partial defense, since it may be possible to infer the selection outcomes of packets by correlating a crafted Packet Stream with the occurrence (not the content) of packets in the export stream that it generates. The rate at which such knowledge could be acquired is limited by the temporal resolution at which reports can be associated with packets, e.g., due to processing and propagation variability, and difficulty in distinguishing report on attack packets from those of background traffic, if present. The association between packets and their reports on which this depends could be removed by padding Export Packets to a constant length and sending them at a constant rate.

We now turn to attacks that can exploit knowledge of selection outcomes. First, with a non-cryptographic Hash Function, knowledge of selection outcomes for a trial stream may be used to further craft a packet set with the common selection property. This has been demonstrated for the modular hash $f(x) = a x + b \text{ mod } k$, for private parameters a , b , and k . With Sampling rate p , knowledge of the Sampling outcomes of roughly $2/p$ is sufficient for the attack to succeed, independent of the values of a , b , and k . With knowledge of the selection outcomes of a larger number of packets, the parameters a , b , and k can be determined; see [GoRe07].

A cryptographic Hash Function employing a private parameter and operating in one of the pseudorandom function modes specified above is not vulnerable to these attacks, even if the selection range is known.

6.2.3.3. Vulnerabilities to Replay Attacks

Since Hash-based Selection is deterministic, any packet or set of packets with known selection properties can be replayed into a network and experience the same selection outcomes provide the Hash Function and its parameters are not changed. Repetition of a single packet may be noticeable to other measurement methods if employed (e.g., collection of flow statistics), whereas a set of distinct packets that appears statistically similar to regular traffic may be less noticeable.

Replay attacks may be mitigated by repeated changing of Hash Function parameters. This also prevents attacks that exploit knowledge of Sampling outcomes, at least if the parameters are changed at least as fast as the knowledge can be acquired by an attacker. In order to preserve the ability to perform trajectory Sampling, parameter change would have to be simultaneous (or approximately so) across all Observation Points.

6.2.4. Choice of Hash Function

The specific choice of Hash Function represents a trade-off between complexity and ease of implementation. Ideally, a cryptographically strong Hash Function employing a private parameter and operating in pseudorandom function mode as specified above would be used, yielding a good emulation of a random packet selection at a target Sampling rate, and giving maximal robustness against the attacks described in the previous section. Unfortunately, there is currently no single Hash Function that fulfills all the requirements.

As detailed in Section 6.2.3, only cryptographic Hash Functions employing a private parameter operating in pseudorandom function mode are sufficiently strong to withstand the range of conceivable attacks. For example, fixed- or variable-length inputs could be hashed using a block cipher (like Advanced Encryption Standard (AES)) in cipher-block-chaining mode. Fixed-length inputs could also be hashed using an iterated cryptographic Hash Function (like MD5 or SHA1), with a private initial vector. For variable-length inputs, an iterated cryptographic Hash Function (like MD5 or SHA1) should employ private string post-pended to the data in addition to a private initial vector. For more details, see the "append-cascade" construction of [BeCK96]. We encourage the use of such cryptographically strong Hash Functions wherever possible.

However, a problem with using such functions is the low performance. As shown for instance in [Henk08], the computation times for MD5 and SHA are about 7-10 times higher compared to non-cryptographic functions. The difference increases for small hash input lengths.

Therefore, it is not assumed that all PSAMP Devices will be capable of applying a cryptographically strong Hash Function to every packet at line rate. For this reason, the Hash Functions listed in this section will be of a weaker variety. Future protocol extensions that employ stronger Hash Functions are highly welcome.

Comparisons of Hash Functions for packet selection and packet digesting with regard to various criteria can be found in [MoND05] and [Henk08].

6.2.4.1. Hash Functions for Packet Selection

If Hash-based packet Selection is applied, the BOB function MUST be used for packet selection operations in order to be compliant with PSAMP. The specification of BOB is given in the appendix. Both the parameter (the init value) and the selection range should be kept private. The initial vector of the Hash Function MUST be configurable out of band to prevent security breaches like exposure of the initial vector content.

Other functions, such as CRC-32 and IPSX, MAY be used. The IPSX function is described in the appendix, and the CRC-32 function is described in [RFC1141]. If CRC-32 is used, the input should first be post-pended with a private string that acts as a parameter, and the modulus of the CRC should also be kept private.

IPSX is simple to implement and was correspondingly about an order of magnitude faster to execute per packet than BOB or CRC-32 [MoND05].

All three Hash Functions evaluated showed relatively poor uniformity with 16-byte input that was drawn from only invariant fields in the IP and TCP/UDP headers (i.e., header fields that do not change from hop to hop). IPSX is inherently limited to 16 bytes.

BOB and CRC-32 exhibit noticeably better uniformity when 4 or more bytes from the payload are also included in the input [MoND05]. Also with other criteria BOB performed quite well [Henk08].

Although the characteristics have been checked for different traffic traces, results cannot be generalized to arbitrary traffic. Since Hash-based Selection is a deterministic function on the Packet Content, it can always be biased towards packets with specific attributes. Furthermore, it should be noted that all Hash Functions were evaluated only for IPv4.

None of these Hash Functions is recommended for cryptographic purposes. Please also note that the use of a private parameter only slightly reduces the vulnerabilities against attacks. As shown in Section 6.2.3, functions that are not cryptographically strong (e.g., BOB and CRC) cannot prevent attackers from crafting packets that are disproportionally selected even if a private parameter is used and the selection range is kept secret.

As described in Section 6.2.2, the input bytes for the Hash Function need to be invariant along the path the packet is traveling. Only with this it is ensured that the same packets are selected at

different Observation Points. Furthermore, they should have a high variability between different packets to generate a high variation in the Hash Range. An evaluation of the variability of different packet header fields can be found in [DuGr00], [HeSZ08], and [Henk08].

If a Hash-based Selection with the BOB function is used with IPv4 traffic, the following input bytes MUST be used.

- IP identification field
- Flags field
- Fragment offset
- Source IP address
- Destination IP address
- A configurable number of bytes from the IP payload, starting at a configurable offset

Due to the lack of suitable IPv6 packet traces, all candidate Hash Functions in [DuGr00], [MoND05], and [Henk08] were evaluated only for IPv4. Due to the IPv6 header fields and address structure, it is expected that there is less randomness in IPv6 packet headers than in IPv4 headers. Nevertheless, the randomness of IPv6 traffic has not yet been evaluated sufficiently to get any evidence. In addition to this, IPv6 traffic profiles may change significantly in the future when IPv6 is used by a broader community.

If a Hash-based Selection with the BOB function is used with IPv6 traffic, the following input bytes MUST be used.

- Payload length (2 bytes)
- Byte number 10,11,14,15,16 of the IPv6 source address
- Byte number 10,11,14,15,16 of the IPv6 destination address
- A configurable number of bytes from the IP payload, starting at a configurable offset. It is recommended to use at least 4 bytes from the IP payload.

The payload itself is not changing during the path. Even if some routers process some extension headers, they are not going to strip them from the packet. Therefore, the payload length is invariant along the path. Furthermore, it usually differs for different packets. The IPv6 address has 16 bytes. The first part is the

network part and contains low variation. The second part is the host part and contains higher variation. Therefore, the second part of the address is used. Nevertheless, the uniformity has not been checked for IPv6 traffic.

6.2.4.2. Hash Functions Suitable for Packet Digesting

For this purpose also the BOB function SHOULD be used. Other functions (such as CRC-32) MAY be used. Among the functions capable of operating with variable-length input, BOB and CRC-32 have the fastest execution, BOB being slightly faster. IPSX is not recommended for digesting because it has a significantly higher collision rate and takes only a fixed-length input.

7. Parameters for the Description of Selection Techniques

This section gives an overview of different alternative selection schemes and their required parameters. In order to be compliant with PSAMP, at least one of proposed schemes MUST be implemented.

The decision whether or not to select a packet is based on a function that is performed when the packet arrives at the selection process. Packet selection schemes differ in the input parameters for the selection process and the functions they require to do the packet selection. The following table gives an overview.

| Scheme | Input parameters | Functions |
|------------------------------|---|--|
| systematic count-based | packet position Sampling pattern | packet counter |
| systematic time-based | arrival time Sampling pattern | clock or timer |
| random n-out-of-N | packet position Sampling pattern (random number list) | packet counter, random numbers |
| uniform probabilistic | Sampling probability | random function |
| non-uniform probabilistic | e.g., packet position, Packet Content(parts) | selection function, probability calc. |
| non-uniform flow-state | e.g., flow state, Packet Content(parts) | selection function, probability calc. |
| property match | Packet Content(parts) or router state | filter function or state discovery |
| hash-based | Packet Content(parts) | Hash Function |

7.1. Description of Sampling Techniques

In this section, we define what elements are needed to describe the most common Sampling techniques. Here the selection function is predefined and given by the Selector ID.

Sampler Description:

```
SELECTOR_ID
SELECTOR_TYPE
SELECTOR_PARAMETERS
```

Where:

SELECTOR_ID:

Unique ID for the packet sampler.

SELECTOR_TYPE:

For Sampling processes, the SELECTOR TYPE defines what Sampling algorithm is used.

Values: Systematic count-based | Systematic time-based | Random
|n-out-of-N | uniform probabilistic | Non-uniform probabilistic |
Non-uniform flow state

SELECTOR_PARAMETERS:

For Sampling processes, the SELECTOR PARAMETERS define the input parameters for the process. Interval length in systematic Sampling means that all packets that arrive in this interval are selected. The spacing parameter defines the spacing in time or number of packets between the end of one Sampling interval and the start of the next succeeding interval.

Case n-out-of-N:

- Population Size N, Sample size n

Case systematic time-based:

- Interval length (in usec), Spacing (in usec)

Case systematic count-based:

- Interval length (in packets), Spacing (in packets)

Case uniform probabilistic (with equal probability per packet):

- Sampling probability p

Case non-uniform probabilistic:

- Calculation function for Sampling probability p (see also Section 5.2.2.4)

Case flow state:

- Information reported for flow state Sampling is not defined in this document (see also Section 5.2.2.4)

7.2. Description of Filtering Techniques

In this section, we define what elements are needed to describe the most common Filtering techniques. The structure closely parallels the one presented for the Sampling techniques.

Filter Description:

SELECTOR_ID
SELECTOR_TYPE
SELECTOR_PARAMETERS

Where:

SELECTOR_ID:

Unique ID for the packet filter. The ID can be calculated under consideration of the SELECTION SEQUENCE and a local ID.

SELECTOR_TYPE:

For Filtering processes, the SELECTOR TYPE defines what Filtering type is used.

Values: Matching | Hashing | Router_state

SELECTOR_PARAMETERS:

For Filtering processes, the SELECTOR PARAMETERS define formally the common property of the packet being filtered. For the filters of type matching and hashing, the definitions have a lot of points in common.

Values:

Case matching:

- Information Element (from [RFC5102])
- Value (type in accordance to [RFC5102])

In case of multiple match criteria, multiple "case matching" has to be bound by a logical AND.

Case hashing:

- Hash Domain (input bits from packet)
 - <Header type = IPv4>
 - <Input bit specification, header part>
 - <Header type = IPv6>
 - <Input bit specification, header part>
 - <payload byte number N>
 - <Input bit specification, payload part>
- Hash Function
 - Hash Function name
 - Length of input key (eliminate 0x bytes)
 - Output value (length M and bitmask)
 - Hash Selection Range, as a list of non-overlapping intervals [start value, end value] where value is in $[0, 2^M - 1]$
 - Additional parameters are dependent on specific Hash Function (e.g., hash input bits (seed))

Notes to input bits for case hashing:

- Input bits can be from header part only, from the payload part only, or from both.

- The bit specification, for the header part, can be specified for IPv4 or IPv6 only, or both.
- In case of IPv4, the bit specification is a sequence of 20 hexadecimal numbers [00,FF] specifying a 20-byte bitmask to be applied to the header.
- In case of IPv6, it is a sequence of 40 hexadecimal numbers [00,FF] specifying a 40-byte bitmask to be applied to the header.
- The bit specification, for the payload part, is a sequence of hexadecimal numbers [00,FF] specifying the bitmask to be applied to the first N bytes of the payload, as specified by the previous field. In case the hexadecimal number sequence is longer than N, only the first N numbers are considered.
- In case the payload is shorter than N, the Hash Function cannot be applied. Other options, like padding with zeros, may be considered in the future.
- A Hash Function cannot be defined on the options field of the IPv4 header, neither on stacked headers of IPv6.
- The Hash Selection Range defines a range of hash values (out of all possible results of the hash operation). If the hash result for a specific packet falls in this range, the packet is selected. If the value is outside the range, the packet is not selected. For example, if the selection interval specification is [1:3], [6:9] all packets are selected for which the hash result is 1,2,3,6,7,8, or 9. In all other cases, the packet is not selected.

Case router state:

- Ingress interface at which the packet arrives equals a specified value
- Egress interface to which the packet is routed equals a specified value
- Packet violated Access Control List (ACL) on the router
- Reverse Path Forwarding (RPF) failed for the packet
- Resource Reservation is insufficient for the packet
- No route is found for the packet
- Origin AS equals a specified value or lies within a given range

- Destination AS equals a specified value or lies within a given range

Note to case router state:

- All router state entries can be linked by AND operators

8. Composite Techniques

Composite schemes are realized by combining the Selector IDs into a Selection Sequence. The Selection Sequence contains all Selector IDs that are applied to the Packet Stream subsequently. Some examples of composite schemes are reported below.

8.1. Cascaded Filtering->Sampling or Sampling->Filtering

If a filter precedes a Sampling process, the role of Filtering is to create a set of "parent populations" from a single stream that can then be fed independently to different Sampling functions, with different parameters tuned for the Population itself (e.g., if streams of different intensity result from Filtering, it may be good to have different Sampling rates). If Filtering follows a Sampling process, the same Selection Fraction and type are applied to the whole stream, independently of the relative size of the streams resulting from the Filtering function. Moreover, also packets not destined to be selected in the Filtering operation will "load" the Sampling function. So, in principle, Filtering before Sampling allows a more accurate tuning of the Sampling procedure, but if filters are too complex to work at full line rate (e.g., because they have to access router state information), Sampling before Filtering may be a need.

8.2. Stratified Sampling

Stratified Sampling is one example for using a composite technique. The basic idea behind stratified Sampling is to increase the estimation accuracy by using a priori information about correlations of the investigated characteristic with some other characteristic that is easier to obtain. The a priori information is used to perform an intelligent grouping of the elements of the parent Population. In this manner, a higher estimation accuracy can be achieved with the same sample size or the sample size can be reduced without reducing the estimation accuracy.

Stratified Sampling divides the Sampling process into multiple steps. First, the elements of the parent Population are grouped into subsets in accordance to a given characteristic. This grouping can be done in multiple steps. Then samples are taken from each subset.

The stronger the correlation between the characteristic used to divide the parent Population (stratification variable) and the characteristic of interest (for which an estimate is sought after), the easier is the consecutive Sampling process and the higher is the stratification gain. For instance, if the dividing characteristic were equal to the investigated characteristic, each element of the subgroup would be a perfect representative of that characteristic. In this case, it would be sufficient to take one arbitrary element out of each subgroup to get the actual distribution of the characteristic in the parent Population. Therefore, stratified Sampling can reduce the costs for the Sampling process (i.e., the number of samples needed to achieve a given level of confidence).

For stratified Sampling, one has to specify classification rules for grouping the elements into subgroups and the Sampling scheme that is used within the subgroups. The classification rules can be expressed by multiple filters. For the Sampling scheme within the subgroups, the parameters have to be specified as described above. The use of stratified Sampling methods for measurement purposes is described for instance in [ClPB93] and [Zseb03].

9. Security Considerations

Security considerations concerning the choice of a Hash Function for Hash-based Selection have been discussed in Section 6.2.3. That section discussed a number of potential attacks to craft Packet Streams that are disproportionately detected and/or discover the Hash Function parameters, the vulnerabilities of different Hash Functions to these attacks, and practices to minimize these vulnerabilities.

In addition to this, a user can gain knowledge about the start and stop triggers in time-based systematic Sampling, e.g., by sending test packets. This knowledge might allow users to modify their send schedule in a way that their packets are disproportionately selected or not selected [GoRe07].

For random Sampling, a cryptographically strong random number generator should be used in order to prevent that an adversary can predict the selection decision [GoRe07].

Further security threats can occur when Sampling parameters are configured or communicated to other entities. The configuration and reporting of Sampling parameters are out of scope of this document. Therefore, the security threats that originate from this kind of communication cannot be assessed with the information given in this document.

Some of these threats can probably be addressed by keeping configuration information confidential and by authenticating entities that configure Sampling. Nevertheless, a full analysis and assessment of threats for configuration and reporting has to be done if configuration or reporting methods are proposed.

10. Contributors

Sharon Goldberg contributed to the security considerations for Hash-based Selection.

Sharon Goldberg
Department of Electrical Engineering
Princeton University
F210-K EQuad
Princeton, NJ 08544,
USA
EMail: goldbe@princeton.edu

11. Acknowledgments

We would like to thank the PSAMP group, especially Benoit Claise and Stewart Bryant, for fruitful discussions and for proofreading the document. We thank Sharon Goldberg for her input on security issues concerning Hash-based Selection.

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.

12.2. Informative References

[AmCa89] Paul D. Amer, Lillian N. Cassel, "Management of Sampled Real-Time Network Measurements", 14th Conference on Local Computer Networks, October 1989, Minneapolis, pages 62-68, IEEE, 1989.

[BeCK96] M. Bellare, R. Canetti and H. Krawczyk, "Pseudorandom Functions Revisited: The Cascade Construction and its Concrete Security", Symposium on Foundations of Computer Science, 1996.

- [ClPB93] K.C. Claffy, George C. Polyzos, Hans-Werner Braun, "Application of Sampling Methodologies to Network Traffic Characterization", Proceedings of ACM SIGCOMM'93, San Francisco, CA, USA, September 13 - 17, 1993.
- [DuGG02] N.G. Duffield, A. Gerber, M. Grossglauser, "Trajectory Engine: A Backend for Trajectory Sampling", IEEE Network Operations and Management Symposium 2002, Florence, Italy, April 15-19, 2002.
- [DuGr00] N.G. Duffield, M. Grossglauser, "Trajectory Sampling for Direct Traffic Observation", Proceedings of ACM SIGCOMM 2000, Stockholm, Sweden, August 28 - September 1, 2000.
- [DuGr04] N.G. Duffield and M. Grossglauser "Trajectory Sampling with Unreliable Reporting", Proc IEEE Infocom 2004, Hong Kong, March 2004.
- [DuLT01] N.G. Duffield, C. Lund, and M. Thorup, "Charging from Sampled Network Usage", ACM Internet Measurement Workshop IMW 2001, San Francisco, USA, November 1-2, 2001.
- [EsVa01] C. Estan and G. Varghese, "New Directions in Traffic Measurement and Accounting", ACM SIGCOMM Internet Measurement Workshop 2001, San Francisco (CA) Nov. 2001.
- [GoRe07] S. Goldberg, J. Rexford, "Security Vulnerabilities and Solutions for Packet Sampling", IEEE Sarnoff Symposium, Princeton, NJ, May 2007.
- [HT52] D.G. Horvitz and D.J. Thompson, "A Generalization of Sampling without replacement from a Finite Universe" J. Amer. Statist. Assoc. Vol. 47, pp. 663-685, 1952.
- [Henk08] Christian Henke, Evaluation of Hash Functions for Multipoint Sampling in IP Networks, Diploma Thesis, TU Berlin, April 2008.
- [HeSZ08] Christian Henke, Carsten Schmoll, Tanja Zseby, Evaluation of Header Field Entropy for Hash-Based Packet Selection, Proceedings of Passive and Active Measurement Conference PAM 2008, Cleveland, Ohio, USA, April 2008.
- [Jenk97] B. Jenkins, "Algorithm Alley", Dr. Dobb's Journal, September 1997.
<http://burtleburtle.net/bob/hash/doobs.html>.

- [JePP92] Jonathan Jedwab, Peter Phaal, Bob Pinna, "Traffic Estimation for the Largest Sources on a Network, Using Packet Sampling with Limited Storage", HP technical report, Management, Mathematics and Security Department, HP Laboratories, Bristol, March 1992, <http://www.hpl.hp.com/techreports/92/HPL-92-35.html>.
- [Moli03] M. Molina, "A scalable and efficient methodology for flow monitoring in the Internet", International Teletraffic Congress (ITC-18), Berlin, Sep. 2003.
- [MoND05] M. Molina, S. Niccolini, N.G. Duffield, "A Comparative Experimental Study of Hash Functions Applied to Packet Sampling", International Teletraffic Congress (ITC-19), Beijing, August 2005.
- [RFC1141] Mallory, T. and A. Kullberg, "Incremental updating of the Internet checksum", RFC 1141, January 1990.
- [RFC1624] Rijsinghani, A., Ed., "Computation of the Internet Checksum via Incremental Update", RFC 1624, May 1994.
- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [RFC3917] Quittek, J., Zseby, T., Claise, B., and S. Zander, "Requirements for IP Flow Information Export (IPFIX)", RFC 3917, October 2004.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [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.
- [RFC5474] Duffield, N., Ed., "A Framework for Packet Selection and Reporting", RFC 5474, March 2009.

- [RFC5476] Claise, B., Ed., "Packet Sampling (PSAMP) Protocol Specifications", RFC 5476, March 2009.
- [RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", RFC 5477, March 2009.
- [Zseb03] T. Zseby, "Stratification Strategies for Sampling-based Non-intrusive Measurement of One-way Delay", Proceedings of Passive and Active Measurement Workshop (PAM 2003), La Jolla, CA, USA, pp. 171-179, April 2003.
- [ZsZC01] Tanja Zseby, Sebastian Zander, Georg Carle. Evaluation of Building Blocks for Passive One-way-delay Measurements. Proceedings of Passive and Active Measurement Workshop (PAM 2001), Amsterdam, The Netherlands, April 23-24, 2001.

Appendix A. Hash Functions

A.1. IP Shift-XOR (IPSX) Hash Function

The IPSX Hash Function is tailored for acting on IP version 4 packets. It exploits the structure of IP packets and in particular the variability expected to be exhibited within different fields of the IP packet in order to furnish a hash value with little apparent correlation with individual packet fields. Fields from the IPv4 and TCP/UDP headers are used as input. The IPSX Hash Function uses a small number of simple instructions.

Input parameters: None

Built-in parameters: None

Output: The output of the IPSX is a 16-bit number

Functioning:

The functioning can be divided into two parts: input selection, whose forms are composite input from various portions of the IP packet, followed by computation of the hash on the composite.

Input Selection:

The raw input is drawn from the first 20 bytes of the IP packet header and the first 8 bytes of the IP payload. If IP options are not used, the IP header has 20 bytes, and hence the two portions adjoin and comprise the first 28 bytes of the IP packet. We now use the raw input as four 32-bit subportions of these 28 bytes. We specify the input by bit offsets from the start of IP header or payload.

f1 = bits 32 to 63 of the IP header, comprising the IP identification field, flags, and fragment offset.

f2 = bits 96 to 127 of the IP header, the source IP address.

f3 = bits 128 to 159 of the IP header, the destination IP address.

f4 = bits 32 to 63 of the IP payload. For a TCP packet, f4 comprises the TCP sequence number followed by the message length. For a UDP packet, f4 comprises the UDP checksum.

Hash Computation:

The hash is computed from f_1 , f_2 , f_3 , and f_4 by a combination of XOR (^), right shift (>>), and left shift (<<) operations. The intermediate quantities h_1 , v_1 , and v_2 are 32-bit numbers.

1. $v_1 = f_1 \wedge f_2;$
2. $v_2 = f_3 \wedge f_4;$
3. $h_1 = v_1 \ll 8;$
4. $h_1 \wedge= v_1 \gg 4;$
5. $h_1 \wedge= v_1 \gg 12;$
6. $h_1 \wedge= v_1 \gg 16;$
7. $h_1 \wedge= v_2 \ll 6;$
8. $h_1 \wedge= v_2 \ll 10;$
9. $h_1 \wedge= v_2 \ll 14;$
10. $h_1 \wedge= v_2 \gg 7;$

The output of the hash is the least significant 16 bits of h_1 .

A.2. BOB Hash Function

The BOB Hash Function is a Hash Function designed for having each bit of the input affecting every bit of the return value and using both 1-bit and 2-bit deltas to achieve the so-called avalanche effect [Jenk97]. This function was originally built for hash table lookup with fast software implementation.

Input parameters:

The input parameters of such a function are:

- the length of the input string (key) to be hashed, in bytes. The elementary input blocks of BOB hash are the single bytes; therefore, no padding is needed.
- an init value (an arbitrary 32-bit number).

Built-in parameters:

The BOB hash uses the following built-in parameter:

- the golden ratio (an arbitrary 32-bit number used in the Hash Function computation: its purpose is to avoid mapping all zeros to all zeros).

Note: The mix sub-function (see mix (a,b,c) macro in the reference code below) has a number of parameters governing the shifts in the registers. The one presented is not the only possible choice.

It is an open point whether these may be considered additional built-in parameters to specify at function configuration.

Output:

The output of the BOB function is a 32-bit number. It should be specified:

- A 32-bit mask to apply to the output
- The Selection Range as a list of non-overlapping intervals [start value, end value] where value is in $[0, 2^{32}]$

Functioning:

The hash value is obtained computing first an initialization of an internal state (composed of three 32-bit numbers, called a, b, c in the reference code below), then, for each input byte of the key the internal state is combined by addition and mixed using the mix sub-function. Finally, the internal state mixed one last time and the third number of the state (c) is chosen as the return value.

```
typedef unsigned long int  ub4;  /* unsigned 4-byte quantities
*/
typedef unsigned          char ub1; /* unsigned 1-byte quantities
*/
```

```
#define hashsize(n) ((ub4)1<<(n))
#define hashmask(n) (hashsize(n)-1)
```

```
/* -----
mix -- mix three 32-bit values reversibly.
```

For every delta with one or two bits set, and the deltas of all three high bits or all three low bits, whether the original value of a,b,c is almost all zero or is uniformly distributed,
 * If mix() is run forward or backward, at least 32 bits in a,b,c have at least 1/4 probability of changing.

* If mix() is run forward, every bit of c will change between 1/3 and 2/3 of the time (well, 22/100 and 78/100 for some 2-bit deltas) mix() was built out of 36 single-cycle latency instructions in a structure that could support 2x parallelism, like so:

```

a -= b;
a -= c; x = (c>>13);
b -= c; a ^= x;
b -= a; x = (a<<8);
c -= a; b ^= x;
c -= b; x = (b>>13);
...

```

Unfortunately, superscalar Pentiums and Sparcs can't take advantage of that parallelism. They've also turned some of those single-cycle latency instructions into multi-cycle latency instructions

-----*/

```

#define mix(a,b,c) \
{ \
  a -= b; a -= c; a ^= (c>>13); \
  b -= c; b -= a; b ^= (a<<8); \
  c -= a; c -= b; c ^= (b>>13); \
  a -= b; a -= c; a ^= (c>>12); \
  b -= c; b -= a; b ^= (a<<16); \
  c -= a; c -= b; c ^= (b>>5); \
  a -= b; a -= c; a ^= (c>>3); \
  b -= c; b -= a; b ^= (a<<10); \
  c -= a; c -= b; c ^= (b>>15); \
}

```

/* -----
hash() -- hash a variable-length key into a 32-bit value
k : the key (the unaligned variable-length array of bytes)
len : the length of the key, counting by bytes
initval : can be any 4-byte value
Returns a 32-bit value. Every bit of the key affects every bit of the return value. Every 1-bit and 2-bit delta achieves avalanche. About 6*len+35 instructions.

The best hash table sizes are powers of 2. There is no need to do mod a prime (mod is so slow!). If you need less than 32 bits, use a bitmask. For example, if you need only 10 bits, do `h = (h & hashmask(10))`, in which case, the hash table should have `hashsize(10)` elements.

If you are hashing `n` strings (`ubl **`)`k`, do it like this: `for (i=0, h=0; i<n; ++i) h = hash(k[i], len[i], h);`

By Bob Jenkins, 1996. bob_jenkins@burtleburtle.net. You may use this code any way you wish, private, educational, or commercial. It's free. See <http://burtleburtle.net/bob/hash/evahash.html>. Use for hash table lookup, or anything where one collision in 2^{32} is acceptable. Do NOT use for cryptographic purposes.

```

----- */
ub4 bob_hash(k, length, initval)
register ub1 *k;      /* the key */
register ub4 length; /* the length of the key */
register ub4 initval; /* an arbitrary value */
{
    register ub4 a,b,c,len;

    /* Set up the internal state */
    len = length;
    a = b = 0x9e3779b9; /*the golden ratio; an arbitrary value
*/
    c = initval;      /* another arbitrary value */

    /*----- handle most of the key */

    while (len >= 12)
    {
        a += (k[0] +((ub4)k[1]<<8) +((ub4)k[2]<<16)
+((ub4)k[3]<<24));
        b += (k[4] +((ub4)k[5]<<8) +((ub4)k[6]<<16)
+((ub4)k[7]<<24));
        c += (k[8] +((ub4)k[9]<<8)
+((ub4)k[10]<<16)+((ub4)k[11]<<24));
        mix(a,b,c);
        k += 12; len -= 12;
    }

    /*----- handle the last 11 bytes */
    c += length;
    switch(len)      /* all the case statements fall through*/
    {
    case 11: c+=((ub4)k[10]<<24);
    case 10: c+=((ub4)k[9]<<16);
    case 9 : c+=((ub4)k[8]<<8);
        /* the first byte of c is reserved for the length */
    case 8 : b+=((ub4)k[7]<<24);
    case 7 : b+=((ub4)k[6]<<16);
    case 6 : b+=((ub4)k[5]<<8);
    case 5 : b+=k[4];
    case 4 : a+=((ub4)k[3]<<24);
    case 3 : a+=((ub4)k[2]<<16);
    }
}

```

```
case 2 : a+=((ub4)k[1]<<8);
case 1 : a+=k[0];
/* case 0: nothing left to add */
}
mix(a,b,c);
/*----- report the result */
return c;
}
```

Authors' Addresses

Tanja Zseby
Fraunhofer Institute for Open Communication Systems
Kaiserin-Augusta-Allee 31
10589 Berlin
Germany
Phone: +49-30-34 63 7153
EMail: tanja.zseby@fokus.fraunhofer.de

Maurizio Molina
DANTE
City House
126-130 Hills Road
Cambridge CB21PQ
United Kingdom
Phone: +44 1223 371 300
EMail: maurizio.molina@dante.org.uk

Nick Duffield
AT&T Labs - Research
Room B-139
180 Park Ave
Florham Park, NJ 07932
USA
Phone: +1 973-360-8726
EMail: duffield@research.att.com

Saverio Niccolini
Network Laboratories, NEC Europe Ltd.
Kurfuerstenanlage 36
69115 Heidelberg
Germany
Phone: +49-6221-9051118
EMail: saverio.niccolini@netlab.nec.de

Frederic Raspall
EPSC-UPC
Dept. of Telematics
Av. del Canal Olympic, s/n
Edifici C4
E-08860 Castelldefels, Barcelona
Spain
EMail: fredri@entel.upc.es