

Internet Engineering Task Force (IETF)
Request for Comments: 6645
Category: Informational
ISSN: 2070-1721

J. Novak
Cisco Systems, Inc.
July 2012

IP Flow Information Accounting and Export Benchmarking Methodology

Abstract

This document provides a methodology and framework for quantifying the performance impact of the monitoring of IP flows on a network device and the export of this information to a Collector. It identifies the rate at which the IP flows are created, expired, and successfully exported as a new performance metric in combination with traditional throughput. The metric is only applicable to the devices compliant with RFC 5470, "Architecture for IP Flow Information Export". The methodology quantifies the impact of the IP flow monitoring process on the network equipment.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see 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/rfc6645>.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
2. Terminology	5
2.1. Existing Terminology	5
2.2. New Terminology	6
3. Flow Monitoring Performance Benchmark	8
3.1. Definition	8
3.2. Device Applicability	8
3.3. Measurement Concept	8
3.4. The Measurement Procedure Overview	9
4. Measurement Setup	11
4.1. Measurement Topology	11
4.2. Baseline DUT Setup	13
4.3. Flow Monitoring Configuration	13
4.4. Collector	19
4.5. Sampling	19
4.6. Frame Formats	19
4.7. Frame Sizes	20
4.8. Flow Export Data Packet Sizes	20
4.9. Illustrative Test Setup Examples	20
5. Flow Monitoring Throughput Measurement Methodology	22
5.1. Flow Monitoring Configuration	23
5.2. Traffic Configuration	24
5.3. Cache Population	25
5.4. Measurement Time Interval	25
5.5. Flow Export Rate Measurement	26
5.6. The Measurement Procedure	27
6. RFC 2544 Measurements	28
6.1. Flow Monitoring Configuration.....	28
6.2. Measurements with the Flow Monitoring Throughput Setup	29
6.3. Measurements with Fixed Flow Export Rate.....	29
7. Flow Monitoring Accuracy	30
8. Evaluating Flow Monitoring Applicability	31
9. Acknowledgements	32
10. Security Considerations	32
11. References	33
11.1. Normative References	33
11.2. Informative References	33
Appendix A. Recommended Report Format	35
Appendix B. Miscellaneous Tests	36
B.1. DUT Under Traffic Load	36
B.2. In-Band Flow Export	36
B.3. Variable Packet Rate	37
B.4. Bursty Traffic	37
B.5. Various Flow Monitoring Configurations	38
B.6. Tests with Bidirectional Traffic	38
B.7. Instantaneous Flow Export Rate	39

1. Introduction

Monitoring IP flows (Flow monitoring) is defined in the "Architecture for IP Flow Information Export" [RFC5470] and related IPFIX documents specified in Section 1.2 of [RFC5470]. It analyzes the traffic using predefined fields from the packet header as keys and stores the traffic and other internal information in the DUT (Device Under Test) memory. This cached flow information is then formatted into records (see Section 2.1 for term definitions) and exported from the DUT to an external data collector for analysis. More details on the measurement architecture are provided in Section 3.3.

Flow monitoring on network devices is widely deployed and has numerous uses in both service-provider and enterprise segments as detailed in the "Requirements for IP Flow Information Export (IPFIX)" [RFC3917]. This document provides a methodology for measuring Flow monitoring performance so that network operators have a framework to measure the impact on the network and network equipment.

This document's goal is to provide a series of methodology specifications for the measurement of Flow monitoring performance in a way that is comparable amongst various implementations, platforms, and vendor devices.

Flow monitoring is, in most cases, run on network devices that also forward packets. Therefore, this document also provides the methodology for [RFC2544] measurements in the presence of Flow monitoring. It is applicable to IPv6 and MPLS traffic with their specifics defined in [RFC5180] and [RFC5695], respectively.

This document specifies a methodology to measure the maximum IP Flow Export Rate that a network device can sustain without impacting the Forwarding Plane, without losing any IP flow information and without compromising IP flow accuracy (see Section 7 for details).

[RFC2544], [RFC5180], and [RFC5695] specify benchmarking of network devices forwarding IPv4, IPv6, and MPLS [RFC3031] traffic, respectively. The methodology specified in this document stays the same for any traffic type. The only restriction may be the DUT's lack of support for Flow monitoring of a particular traffic type.

A variety of different DUT architectures exist that are capable of Flow monitoring and export. As such, this document does not attempt to list the various white-box variables (e.g., CPU load, memory utilization, hardware resources utilization, etc.) that could be gathered as they always help in comparison evaluations. A more complete understanding of the stress points of a particular device

can be attained using this internal information, and the tester MAY choose to gather this information during the measurement iterations.

2. Terminology

The terminology used in this document is based on that defined in [RFC5470], [RFC2285], and [RFC1242], as summarized in Section 2.1. The only new terms needed for this methodology are defined in Section 2.2.

Additionally, 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].

2.1. Existing Terminology

Device Under Test (DUT)	[RFC2285, Section 3.1.1]
Flow	[RFC5101, Section 2]
Flow Key	[RFC5101, Section 2]
Flow Record	[RFC5101, Section 2]
Template Record	[RFC5101, Section 2]
Observation Point	[RFC5470, Section 2]
Metering Process	[RFC5470, Section 2]
Exporting Process	[RFC5470, Section 2]
Exporter	[RFC5470, Section 2]
Collector	[RFC5470, Section 2]
Control Information	[RFC5470, Section 2]
Data Stream	[RFC5470, Section 2]
Flow Expiration	[RFC5470, Section 5.1.1]
Flow Export	[RFC5470, Section 5.1.2]
Throughput	[RFC1242, Section 3.17]

2.2. New Terminology

2.2.1. Cache

Definition:

Memory area held and dedicated by the DUT to store Flow information prior to the Flow Expiration.

2.2.2. Cache Size

Definition:

The size of the Cache in terms of how many entries the Cache can hold.

Discussion:

This term is typically represented as a configurable option in the particular Flow monitoring implementation. Its highest value will depend on the memory available in the network device.

Measurement units:

Number of Cache entries

2.2.3. Active Timeout

Definition:

For long-running Flows, the time interval after which the Metering Process expires a Cache entry to ensure Flow data is regularly updated.

Discussion:

This term is typically presented as a configurable option in the particular Flow monitoring implementation. See Section 5.1.1 of [RFC5470] for a more detailed discussion.

Flows are considered long running when they last longer than several multiples of the Active Timeout. If the Active Timeout is zero, then Flows are considered long running if they contain many more packets (tens of packets) than usually observed in a single transaction.

Measurement units:

Seconds

2.2.4. Idle Timeout

Definition:

The time interval used by the Metering Process to expire an entry from the Cache when no more packets belonging to that specific Cache entry have been observed during the interval.

Discussion:

Idle Timeout is typically represented as a configurable option in the particular Flow monitoring implementation. See Section 5.1.1 of [RFC5470] for more detailed discussion. Note that some documents in the industry refer to "Idle Timeout" as "inactive timeout".

Measurement units:

Seconds

2.2.5. Flow Export Rate

Definition:

The number of Cache entries that expire from the Cache (as defined by the Flow Expiration term) and are exported to the Collector within a measurement time interval. There SHOULD NOT be any export filtering, so that all the expired Cache entries are exported. If there is export filtering and it can't be disabled, this MUST be indicated in the measurement report.

The measured Flow Export Rate MUST include both the Data Stream and the Control Information, as defined in Section 2 of [RFC5470].

Discussion:

The Flow Export Rate is measured using Flow Export data observed at the Collector by counting the exported Flow Records during the measurement time interval (see Section 5.4). The value obtained is an average of the instantaneous export rates observed during the measurement time interval. The smallest possible measurement interval (if attempting to measure a nearly instantaneous export rate rather than average export rate on the DUT) is limited by the export capabilities of the particular Flow monitoring implementation (when physical-layer issues between the DUT and the Collector are excluded).

Measurement units:

Number of Flow Records per second

3. Flow Monitoring Performance Benchmark

3.1. Definition

Flow Monitoring Throughput

Definition:

The maximum Flow Export Rate the DUT can sustain without losing a single Cache entry. Additionally, for packet forwarding devices, the maximum Flow Export Rate the DUT can sustain without dropping packets in the Forwarding Plane (see Figure 1).

Measurement units:

Number of Flow Records per second

Discussion:

The losses of Cache entries, or forwarded packets per this definition are assumed to happen due to the lack of DUT resources to process any additional traffic information or lack of resources to process Flow Export data. The physical-layer issues, like insufficient bandwidth from the DUT to the Collector or lack of Collector resources, MUST be excluded as detailed in Section 4.

3.2. Device Applicability

The Flow monitoring performance metric is applicable to network devices that deploy the architecture described in [RFC5470]. These devices can be network packet forwarding devices or appliances that analyze traffic but do not forward traffic (e.g., probes, sniffers, replicators).

This document does not intend to measure Collector performance, it only requires sufficient Collector resources (as specified in Section 4.4) in order to measure the DUT characteristics.

3.3. Measurement Concept

Figure 1 presents the functional block diagram of the DUT. The traffic in the figure represents test traffic sent to the DUT and forwarded by the DUT, if possible. When testing devices that do not act as network packet forwarding devices (such as probes, sniffers, and replicators), the Forwarding Plane is simply an Observation Point as defined in Section 2 of [RFC5470]. The Throughput of such devices will always be zero, and the only applicable performance metric is the Flow Monitoring Throughput. Netflow is specified by [RFC3954].

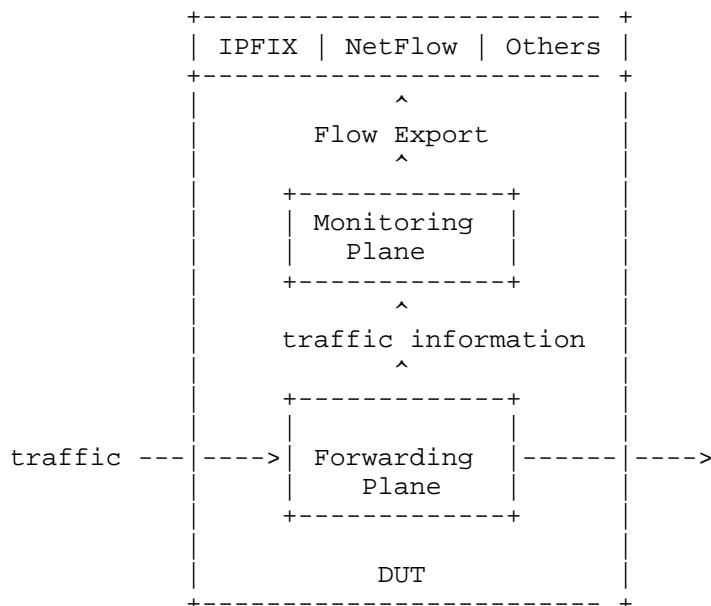


Figure 1. The Functional Block Diagram of the DUT

Flow monitoring is represented in Figure 1 by the Monitoring Plane; it is enabled as specified in Section 4.3. It uses the traffic information provided by the Forwarding Plane and configured Flow Keys to create Cache entries representing the traffic forwarded (or observed) by the DUT in the DUT Cache. The Cache entries are expired from the Cache depending on the Cache configuration (e.g., the Active and Idle Timeouts, the Cache Size), number of Cache entries, and the traffic pattern. The Cache entries are used by the Exporting Process to format the Flow Records, which are then exported from the DUT to the Collector (see Figure 2 in Section 4).

The Forwarding Plane and Monitoring Plane represent two separate functional blocks, each with its own performance capability. The Forwarding Plane handles user data packets and is fully characterized by the metrics defined by [RFC1242].

The Monitoring Plane handles Flows that reflect the analyzed traffic. The metric for Monitoring Plane performance is the Flow Export Rate, and the benchmark is the Flow Monitoring Throughput.

3.4. The Measurement Procedure Overview

The measurement procedure is fully specified in Sections 4, 5, and 6. This section provides an overview of principles for the measurements.

The basic measurement procedure of the performance characteristics of a DUT with Flow monitoring enabled is a conventional Throughput measurement using a search algorithm to determine the maximum packet rate at which none of the offered packets and corresponding Flow Records are dropped by the DUT as described in [RFC1242] and Section 26.1 of [RFC2544].

The DUT with Flow monitoring enabled contains two functional blocks that need to be measured using characteristics applicable to one or both blocks (see Figure 1). See Sections 3.4.1 and 3.4.2 for further discussion.

On one hand, the Monitoring Plane and Forwarding Plane (see Figure 1) need to be looked at as two independent blocks, and the performance of each measured independently. On the other hand, when measuring the performance of one, the status and performance of the other MUST be known and benchmarked when both are present.

3.4.1. Monitoring Plane Performance Measurement

The Flow Monitoring Throughput MUST be (and can only be) measured with one packet per Flow as specified in Section 5. This traffic type represents the most demanding traffic from the Flow monitoring point of view and will exercise the Monitoring Plane (see Figure 1) of the DUT most. In this scenario, every packet seen by the DUT creates a new Cache entry and forces the DUT to fill the Cache instead of just updating the packet and byte counters of an already existing Cache entry.

The exit criteria for the Flow Monitoring Throughput measurement are one of the following (e.g., if any of the conditions are reached):

- a. The Flow Export Rate at which the DUT starts to lose Flow Information or the Flow Information gets corrupted.
- b. The Flow Export Rate at which the Forwarding Plane starts to drop or corrupt packets (if the Forwarding Plane is present).

A corrupted packet here means packet header corruption (resulting in the cyclic redundancy check failure on the transmission level and consequent packet drop) or packet payload corruption, which leads to lost application-level data.

3.4.2. Forwarding Plane Performance Measurement

The Forwarding Plane (see Figure 1) performance metrics are fully specified by [RFC1242] and MUST be measured accordingly. A detailed traffic analysis (see below) with relation to Flow monitoring MUST be

performed prior of any [RFC2544] measurements. Most importantly, the Flow Export Rate caused by the test traffic during an [RFC2544] measurement MUST be known and reported.

The required test traffic analysis mainly involves the following:

- a. Which packet header parameters are incremented or changed during traffic generation.
- b. Which Flow Keys the Flow monitoring configuration uses to generate Flow Records.

The performance metrics described in RFC 1242 can be measured in one of the three modes:

- a. As a baseline of forwarding performance without Flow monitoring.
- b. At a certain level of Flow monitoring activity specified by a Flow Export Rate lower than the Flow Monitoring Throughput.
- c. At the maximum level of Flow monitoring performance, e.g., using traffic conditions representing a measurement of Flow Monitoring Throughput.

The above mentioned measurement mode in point a. represents an ordinary Throughput measurement specified in RFC 2544. The details of how to set up the measurements in points b. and c. are given in Section 6.

4. Measurement Setup

This section concentrates on the setup of all components necessary to perform Flow monitoring performance measurement. The recommended reporting format can be found in Appendix A.

4.1. Measurement Topology

The measurement topology described in this section is applicable only to the measurements with packet forwarding network devices. The possible architectures and implementation of the traffic monitoring appliances (see Section 3.2) are too various to be covered in this document. Instead of the Forwarding Plane, these appliances generally have some kind of feed (e.g., an optical splitter, an interface sniffing traffic on a shared media, or an internal channel on the DUT providing a copy of the traffic) providing the information about the traffic necessary for Flow monitoring analysis. The measurement topology then needs to be adjusted to the appliance architecture and MUST be part of the measurement report.

The measurement setup is identical to that used by [RFC2544], with the addition of a Collector to analyze the Flow Export (see Figure 2).

In the measurement topology with unidirectional traffic, the traffic is transmitted from the sender to the receiver through the DUT. The received traffic is analyzed to check that it is identical to the generated traffic.

The ideal way to implement the measurement is by using a single device to provide the sender and receiver capabilities with one sending port and one receiving port. This allows for an easy check as to whether all the traffic sent by the sender was re-transmitted by the DUT and received at the receiver.




Figure 2. Measurement Topology with Unidirectional Traffic

The DUT's export interface (connecting the Collector) MUST NOT be used for forwarding test traffic but only for the Flow Export data containing the Flow Records. In all measurements, the export interface MUST have enough bandwidth to transmit Flow Export data without congestion. In other words, the export interface MUST NOT be a bottleneck during the measurement.

The traffic receiver MUST have sufficient resources to measure all test traffic transferred successfully by the DUT. This may be checked through measurements with and without the DUT.

Note that more complex topologies might be required. For example, if the effects of enabling Flow monitoring on several interfaces is of concern, or the maximum speed of media transmission is less than the DUT Throughput, the topology can be expanded with several input and output ports. However, the topology MUST be clearly written in the measurement report.

4.2. Baseline DUT Setup

The baseline DUT setup and the way the setup is reported in the measurement results is fully specified in Section 7 of [RFC2544].

The baseline DUT configuration might include other features, like packet filters or quality of service on the input and/or output interfaces, if there is the need to study Flow monitoring in the presence of those features. The Flow monitoring measurement procedures do not change in this case. Consideration needs to be made when evaluating measurement results to take into account the possible change of packet rates offered to the DUT and Flow monitoring after application of the features to the configuration. Any such feature configuration MUST be part of the measurement report.

The DUT export interface (see Figure 2) SHOULD be configured with sufficient output buffers to avoid dropping the Flow Export data due to a simple lack of resources in the interface hardware. The applied configuration MUST be part of the measurement report.

The test designer has the freedom to run tests in multiple configurations. It is therefore possible to run both non-production and real deployment configurations in the laboratory, according to the needs of the tester. All configurations MUST be part of the measurement report.

4.3. Flow Monitoring Configuration

This section covers all of the aspects of the Flow monitoring configuration necessary on the DUT in order to perform the Flow monitoring performance measurement. The necessary configuration has a number of components (see [RFC5470]), namely Observation Points, Metering Process, and Exporting Process as detailed below.

The DUT MUST support the Flow monitoring architecture as specified by [RFC5470]. The DUT SHOULD support IPFIX [RFC5101] to allow a meaningful results comparison due to the standardized export protocol.

The DUT configuration, any existing Cache, and Cache entries MUST be erased before the application of any new configuration for the currently executed measurement.

4.3.1. Observation Points

The Observation Points specify the interfaces and direction in which the Flow monitoring traffic analysis is to be performed.

The (*) in Figure 2 designates the Observation Points in the default configuration. Other DUT Observation Points might be configured depending on the specific measurement needs as follows:

- a. ingress port/ports only
- b. egress port/ports only
- c. both ingress and egress

This test topology corresponds to unidirectional traffic only with traffic analysis performed on the input and/or output interface. Testing with bidirectional traffic is discussed in Appendix B.

Generally, the placement of Observation Points depends upon the position of the DUT in the deployed network and the purpose of Flow monitoring. See [RFC3917] for detailed discussion. The measurement procedures are otherwise the same for all these possible configurations.

In the case of both ingress and egress Flow monitoring being enabled on one DUT, the resulting analysis should consider that each Flow will be represented in the DUT Cache by two Flow Records (one for each direction). Therefore, the Flow Export will also contain those two Flow Records.

If more than one Observation Point for one direction is defined on the DUT, the traffic passing through each of the Observation Points MUST be configured in such a way that it creates Flows and Flow Records that do not overlap. Each packet (or set of packets if measuring more than one packet per Flow - see Section 6.3.1) sent to the DUT on different ports still creates one unique Flow Record.

The specific Observation Points and associated monitoring direction MUST be included as part of the measurement report.

4.3.2. Metering Process

The Metering Process MUST be enabled in order to create the Cache in the DUT and configure the Cache related parameters.

The Cache Size available to the DUT MUST be known and taken into account when designing the measurement as specified in Section 5. Typically, the Cache Size will be present in the "show" commands of the Flow monitoring process, in either the actual configuration or the product documentation from the DUT vendor. The Cache Size MUST have a fixed value for the entire duration of the measurement. This method is not applicable to benchmarking any Flow monitoring applications that dynamically change their Cache Size.

The configuration of the Metering Process MUST be included as part of the measurement report. For example, when a Flow monitoring implementation uses timeouts to expire entries from the Cache, the Cache's Idle and Active Timeouts MUST be known and taken into account when designing the measurement as specified in Section 5. If the Flow monitoring implementation allows only timeouts equal to zero (e.g., immediate timeout or non-existent Cache), then the measurement conditions in Section 5 are fulfilled inherently without any additional configuration. The DUT simply exports information about every packet immediately, subject to the Flow Export Rate definition in Section 2.2.5.

If the Flow monitoring implementation allows configuration of multiple Metering Processes on a single DUT, the exact configuration of each process MUST be included in the measurement report. Only measurements with the same number of Metering Processes can be compared.

The Cache Size and the Idle and Active Timeouts MUST be included in the measurement report.

4.3.3. Exporting Process

The Exporting Process MUST be configured in order to export the Flow Record data to the Collector.

The Exporting Process MUST be configured in such a way that all Flow Records from all configured Observation Points are exported towards the Collector, after the expiration policy, which is composed of the Idle and Active Timeouts and Cache Size.

The Exporting Process SHOULD be configured with IPFIX [RFC5101] as the protocol used to format the Flow Export data. If the Flow monitoring implementation does not support IPFIX, proprietary protocols MAY be used. Only measurements with the same export protocol SHOULD be compared since the protocols may differ in their export efficiency. The export efficiency might also be influenced by the Template Record used and the ordering of the individual export fields within the template.

The Template Records used by the tested implementations SHOULD be analyzed and documented as part of the measurement report. Ideally, only tests with same Template Records should be compared.

Various Flow monitoring implementations might use different default values regarding the export of Control Information [RFC5470]; therefore, the Flow Export corresponding to Control Information SHOULD be analyzed and reported as a separate item on the measurement report. The export of Control Information SHOULD always be configured consistently across all testing and configured to the minimal possible value. Ideally, just one set of Control Information should be exported during each measurement. Note that Control Information includes options and Template Records [RFC5470].

Section 10 of [RFC5101] and Section 8.1 of [RFC5470] discuss the possibility of deploying various transport-layer protocols to deliver Flow Export data from the DUT to the Collector. The selected protocol MUST be included in the measurement report. Only benchmarks with the same transport-layer protocol SHOULD be compared. If the Flow monitoring implementation allows the use of multiple transport-layer protocols, each of the protocols SHOULD be measured in a separate measurement run and the results reported independently in the measurement report.

If a reliable transport protocol is used for the transmission of the Flow Export data from the DUT, the configuration of the Transport session MUST allow for non-blocking data transmission. An example of parameters to look at would be the TCP window size and maximum segment size (MSS). The most substantial transport-layer parameters should be included in the measurement report.

4.3.4. Flow Records

A Flow Record contains information about a specific Flow observed at an Observation Point. A Flow Record contains measured properties of the Flow (e.g., the total number of bytes for all the Flow packets) and usually characteristic properties of the Flow (e.g., source IP address).

The Flow Record definition is implementation specific. A Flow monitoring implementation might allow for only a fixed Flow Record definition, based on the most common IP parameters in the IPv4 or IPv6 headers -- for example, source and destination IP addresses, IP protocol numbers, or transport-level port numbers. Another implementation might allow the user to define their own arbitrary Flow Record to monitor the traffic. The only requirement for the measurements defined in this document is the need for a large

number of Cache entries in the Cache. The Flow Keys needed to achieve that will typically be source and destination IP addresses and transport-level port numbers.

The recommended full IPv4, IPv6, or MPLS Flow Record is shown below. The IP address indicates either IPv4 or IPv6, depending on the traffic type being tested. The Flow Record configuration is Flow monitoring implementation-specific; therefore, the examples below cannot provide an exact specification of individual entries in each Flow Record. The best set of key fields to use is left to the test designer using the capabilities of the specific Flow monitoring implementation.

Flow Keys:

- Source IP address
- Destination IP address
- MPLS label (for MPLS traffic type only)
- Transport-layer source port
- Transport-layer destination port
- IP protocol number (IPv6 next header)
- IP type of service (IPv6 traffic class)

Other fields:

- Packet counter
- Byte counter

Table 1: Recommended Configuration

If the Flow monitoring allows for user-defined Flow Records, the minimal Flow Record configurations allowing large numbers of Cache entries are, for example:

Flow Keys:

- Source IP address
- Destination IP address

Other fields:

- Packet counter

or:

Flow Keys:

- Transport-layer source port
- Transport-layer destination port

Other fields:

- Packet counter

Table 2: User-Defined Configuration

The Flow Record configuration MUST be clearly noted in the measurement report. The Flow Monitoring Throughput measurements on different DUTs, or different Flow monitoring implementations, MUST be only compared for exactly the same Flow Record configuration.

4.3.5. Flow Monitoring with Multiple Configurations

The Flow monitoring architecture as specified in [RFC5470] allows for more complicated configurations with multiple Metering and Exporting Processes on a single DUT. Depending on the particular Flow monitoring implementation, it might affect the measured DUT performance. Therefore, the measurement report should contain information about how many Metering and Exporting Processes were configured on the DUT for the selected Observation Points.

The examples of such possible configurations are:

- a. Several Observation Points with a single Metering Process and a single Exporting Process.
- b. Several Observation Points, each with one Metering Process but all using just one instance of Exporting Process.
- c. Several Observation Points with per-Observation-Point Metering Process and Exporting Process.

4.3.6. MPLS Measurement Specifics

The Flow Record configuration for measurements with MPLS encapsulated traffic SHOULD contain the MPLS label. For this document's purposes, "MPLS Label" is the entire 4 byte MPLS header. Typically, the label of the interest will be at the top of the label stack, but this depends on the details of the MPLS test setup.

The tester SHOULD ensure that the data received by the Collector contains the expected MPLS labels.

The MPLS forwarding performance document [RFC5695] specifies a number of possible MPLS label operations to test. The Observation Points MUST be placed on all the DUT test interfaces where the particular MPLS label operation takes place. The performance measurements SHOULD be performed with only one MPLS label operation at the time.

The DUT MUST be configured in such a way that all the traffic is subject to the measured MPLS label operation.

4.4. Collector

The Collector is needed in order to capture the Flow Export data, which allows the Flow Monitoring Throughput to be measured.

The Collector can be used exclusively as a capture device, providing just hexadecimal format of the Flow Export data. In such a case, it does not need to have any additional Flow Export decoding capabilities and all the decoding is done offline.

However, if the Collector is also used to decode the Flow Export data, it SHOULD support IPFIX [RFC5101] for meaningful results analysis. If proprietary Flow Export is deployed, the Collector MUST support it; otherwise, the Flow Export data analysis is not possible.

The Collector MUST be capable of capturing the export packets sent from the DUT at the full rate without losing any of them. When using reliable transport protocols (see also Section 4.3.3) to transmit Flow Export data, the Collector MUST have sufficient resources to guarantee non-blocking data transmission on the transport-layer session.

During the analysis, the Flow Export data needs to be decoded and the received Flow Records counted.

The capture buffer MUST be cleared at the beginning of each measurement.

4.5. Sampling

Packet sampling and flow sampling is out of the scope of this document. This document applies to situations without packet, flow, or export sampling.

4.6. Frame Formats

Flow monitoring itself is not dependent in any way on the media used on the input and output ports. Any media can be used as supported by the DUT and the test equipment. This applies both to data forwarding interfaces and to the export interface (see Figure 2).

At the time of this writing, the most common transmission media and corresponding frame formats (e.g., Ethernet, Packet over SONET) for IPv4, IPv6, and MPLS traffic are specified within [RFC2544], [RFC5180], and [RFC5695].

The presented frame formats MUST be recorded in the measurement report.

4.7. Frame Sizes

Frame sizes of the traffic to be analyzed by the DUT are specified in Section 9 of [RFC2544] for Ethernet type interfaces (64, 128, 256, 1024, 1280, 1518 bytes) and in Section 5 of [RFC5180] for Packet over SONET interfaces (47, 64, 128, 256, 1024, 1280, 1518, 2048, 4096 bytes).

When measuring with large frame sizes, care needs to be taken to avoid any packet fragmentation on the DUT interfaces that could negatively affect measured performance values.

The presented frame sizes MUST be recorded in the measurement report.

4.8. Flow Export Data Packet Sizes

The Flow monitoring performance will be affected by the packet size that the particular implementation uses to transmit Flow Export data to the Collector. The used packet size MUST be part of the measurement report and only measurements with same packet sizes SHOULD be compared.

The DUT export interface (see Figure 2) maximum transmission unit (MTU) SHOULD be configured to the largest available value for the media. The Flow Export MTU MUST be recorded in the measurement report.

4.9. Illustrative Test Setup Examples

The examples below represent a hypothetical test setup to clarify the use of Flow monitoring parameters and configuration, together with traffic parameters to test Flow monitoring. The actual benchmarking specifications are in Sections 5 and 6.

4.9.1. Example 1 - Idle Timeout Flow Expiration

The traffic generator sends 1000 packets per second in 10000 defined streams, each stream identified by a unique destination IP address. Therefore, each stream has a packet rate of 0.1 packets per second.

The packets are sent in a round-robin fashion (stream 1 to 10000) while incrementing the destination IP address for each sent packet. After a packet for stream 10000 is sent, the next packet destination IP address corresponds to stream 1's address again.

The configured Cache Size is 20000 Flow Records. The configured Active Timeout is 100 seconds, and the Idle Timeout is 5 seconds.

Flow monitoring on the DUT uses the destination IP address as the Flow Key.

A packet with the destination IP address equal to A is sent every 10 seconds, so the Cache entry is refreshed in the Cache every 10 seconds. However, the Idle Timeout is 5 seconds, so the Cache entries will expire from the Cache due to the Idle Timeout, and when a new packet is sent with the same IP address A, it will create a new entry in the Cache. This behavior depends upon the design and efficiency of the Cache ager, and incidences of multi-packet flows observed during this test should be noted.

The measured Flow Export Rate in this case will be 1000 Flow Records per second since every single sent packet will always create a new Cache entry and 1000 packets per second are sent.

The expected number of Cache entries in the Cache during the whole measurement is around 5000. It corresponds to the Idle Timeout being 5 seconds; during those five seconds, 5000 entries are created. This expectation might change in real measurement setups with large Cache Sizes and a high packet rate where the DUT's actual export rate might be limited and lower than the Flow Expiration activity caused by the traffic offered to the DUT. This behavior is entirely implementation-specific.

4.9.2. Example 2 - Active Timeout Flow Expiration

The traffic generator sends 1000 packets per second in 100 defined streams, each stream identified by a unique destination IP address. Each stream has a packet rate of 10 packets per second. The packets are sent in a round-robin fashion (stream 1 to 100) while incrementing the destination IP address for each sent packet. After a packet for stream 100 is sent, the next packet destination IP address corresponds to stream 1's address again.

The configured Cache Size is 1000 Flow Records. The configured Active Timeout is 100 seconds. The Idle Timeout is 10 seconds.

Flow monitoring on the DUT uses the destination IP address as the Flow Key.

After the first 100 packets are sent, 100 Cache entries will have been created in the Flow monitoring Cache. The subsequent packets will be counted against the already created Cache entries since the destination IP address (Flow Key) has already been seen by the DUT (provided the Cache entries did not expire yet as described below).

A packet with the destination IP address equal to A is sent every 0.1 second, so the Cache entry is refreshed in the Cache every 0.1 second, while the Idle Timeout is 10 seconds. In this case, the Cache entries will not expire until the Active Timeout expires, e.g., they will expire every 100 seconds and then the Cache entries will be created again.

If the test measurement time is 50 seconds from the start of the traffic generator, then the measured Flow Export Rate is 0 since during this period nothing expired from the Cache.

If the test measurement time is 100 seconds from the start of the traffic generator, then the measured Flow Export Rate is 1 Flow Record per second.

If the test measurement time is 290 seconds from the start of the traffic generator, then the measured Flow Export Rate is $\frac{2}{3}$ of a Flow Record per second since the Cache expired the same number of Flows twice (100) during the 290-seconds period.

5. Flow Monitoring Throughput Measurement Methodology

Objective:

To measure the Flow monitoring performance in a manner that is comparable between different Flow monitoring implementations.

Metric definition:

Flow Monitoring Throughput - see Section 3.

Discussion:

Different Flow monitoring implementations might choose to handle Flow Export from a partially empty Cache differently than in the case of the Cache being fully occupied. Similarly, software- and hardware-based DUTs can handle the same situation as stated above differently. The purpose of the benchmark measurement in this section is to define one measurement procedure covering all the possible behaviors.

The only criteria is to measure as defined here until Flow Record or packet losses are seen. The decision whether to dive deeper into the conditions under which the packet losses happen is left to the tester.

5.1. Flow Monitoring Configuration

Cache Size

Cache Size configuration is dictated by the expected position of the DUT in the network and by the chosen Flow Keys of the Flow Record. The number of unique sets of Flow Keys that the traffic generator (sender) provides should be multiple times larger than the Cache Size. This ensures that the existing Cache entries are never updated by a packet from the sender before the particular Flow Expiration and Flow Export. This condition is simple to fulfill with linearly incremented Flow Keys (for example, IP addresses or transport-layer ports) where the range of values must be larger than the Cache Size. When randomized traffic generation is in use, the generator must ensure that the same Flow Keys are not repeated within a range of randomly generated values.

The Cache Size **MUST** be known in order to define the measurement circumstances properly. Typically, the Cache Size will be found using the "show" commands of the Flow monitoring implementation in the actual configuration or in the product documentation from the vendor.

Idle Timeout

Idle Timeout is set (if configurable) to the minimum possible value on the DUT. This ensures that the Cache entries are expired as soon as possible and exported out of the DUT Cache. It **MUST** be known in order to define the measurement circumstances completely and equally across implementations.

Active Timeout

Active Timeout is set (if configurable) to a value equal to or higher than the Idle Timeout. It **MUST** be known in order to define the measurement circumstances completely and equally across implementations.

Flow Keys Definition:

The test needs large numbers of unique Cache entries to be created by incrementing values of one or several Flow Keys. The number of unique combinations of Flow Keys values **SHOULD** be several times larger than the DUT Cache Size. This makes sure that any incoming packet will never refresh any already existing Cache entry.

The availability of Cache Size, Idle Timeout, and Active Timeout as configuration parameters is implementation-specific. If the Flow monitoring implementation does not support these parameters, the test possibilities, as specified by this document, are restricted. Some

testing might be viable if the implementation follows the guidance provided in the [IPFIX-CONFIG] document and is considered on a case-by-case basis.

5.2. Traffic Configuration

Traffic Generation

The traffic generator needs to increment the Flow Keys values with each sent packet. This way, each packet represents one Cache entry in the DUT Cache.

A particular Flow monitoring implementation might choose to deploy a hashing mechanism to match incoming data packets to a certain Flow. In such a case, the combination of how the traffic is constructed and the hashing might influence the DUT Flow monitoring performance. For example, if IP addresses are used as Flow Keys, this means there could be a performance difference for linearly incremented addresses (in ascending or descending order) as opposed to IP addresses randomized in a certain range. If randomized IP address sequences are used, then the traffic generator needs to be able to reproduce the randomization (e.g., the same set of IP addresses sent in the same order in different test runs) in order to compare various DUTs and Flow monitoring implementations.

If the test traffic rate is below the maximum media rate for the particular packet size, the traffic generator **MUST** send the packets in equidistant time intervals. Traffic generators that do not fulfill this condition **MUST NOT** and cannot be used for the Flow Monitoring Throughput measurement. An example of this behavior is if the test traffic rate is one half of the media rate. The traffic generator achieves this rate by sending packets each half of each second at the full media rate and sending nothing for the second half of each second. In such conditions, it would be impossible to distinguish if the DUT failed to handle the Flows due to the shortage of input buffers during the burst or due to the limits in the Flow monitoring performance.

Measurement Duration

The measurement duration (e.g., how long the test traffic is sent to the DUT) **MUST** be at least two-times longer than the Idle Timeout; otherwise, no Flow Export would be seen. The measurement duration **SHOULD** guarantee that the number of Cache entries created during the measurement exceeds the available Cache Size.

5.3. Cache Population

The product of the Idle Timeout and the packet rate offered to the DUT (Cache population) during one measurement determines the total number of Cache entries in the DUT Cache during the measurement (while taking into account some margin for dynamic behavior during high DUT loads when processing the Flows).

The Flow monitoring implementation might behave differently depending on the relation of the Cache population to the available Cache Size during the measurement. This behavior is fully implementation-specific and will also be influenced if the DUT architecture is software based or hardware based.

The Cache population (if it is lower or higher than the available Cache Size) during a particular benchmark measurement SHOULD be noted, and mainly only measurements with the same Cache population SHOULD be compared.

5.4. Measurement Time Interval

The measurement time interval is the time value that is used to calculate the measured Flow Export Rate from the captured Flow Export data. It is obtained as specified below.

RFC 2544 specifies, with the precision of the packet beginning and ending, the time intervals to be used to measure the DUT time characteristics. In the case of a Flow Monitoring Throughput measurement, the start and stop time needs to be clearly defined, but the granularity of this definition can be limited to just marking the start and stop time with the start and stop of the traffic generator. This assumes that the traffic generator and DUT are collocated and the variance in transmission delay from the generator to the DUT is negligible as compared to the total time of traffic generation.

The measurement start time:

the time when the traffic generator is started

The measurement stop time: the time when the traffic generator is stopped

The measurement time interval is then calculated as the difference (stop time) - (start time) - (Idle Timeout).

This supposes that the Cache Size is large enough that the time needed to fill it with Cache entries is longer than the Idle Timeout. Otherwise, the time needed to fill the Cache needs to be used to calculate the measurement time interval in place of the Idle Timeout.

Instead of measuring the absolute values of the stop and start times, it is possible to set up the traffic generator to send traffic for a certain predefined time interval, which is then used in the above definition instead of the difference (stop time) - (start time).

The Collector **MUST** stop collecting the Flow Export data at the measurement stop time.

The Idle Timeout (or the time needed to fill the Cache) causes delay of the Flow Export data behind the test traffic that is analyzed by the DUT. For example, if the traffic starts at time point X, Flow Export will start only at the time point X + Idle Timeout (or X + time to fill the Cache). Since Flow Export capture needs to stop with the traffic (because that's when the DUT stops processing the Flows at the given rate), the time interval during which the DUT kept exporting data is shorter by the Idle Timeout than the time interval when the test traffic was sent from the traffic generator to the DUT.

5.5. Flow Export Rate Measurement

The Flow Export Rate needs to be measured in two consequent steps. The purpose of the first step (point a. below) is to gain the actual value for the rate; the second step (point b. below) needs to be done in order to verify that no Flow Record are dropped during the measurement:

- a. In the first step, the captured Flow Export data **MUST** be analyzed only for the capturing interval (measurement time interval) as specified in Section 5.4. During this period, the DUT is forced to process Cache entries at the rate the packets are sent. When traffic generation finishes, the behavior when emptying the Cache is completely implementation-specific; therefore, the Flow Export data from this period cannot be used for benchmarking.
- b. In the second step, all the Flow Export data from the DUT **MUST** be captured in order to determine the Flow Record losses. It needs to be taken into account that especially when large Cache Sizes (in order of magnitude of hundreds of thousands of entries and higher) are in use, the Flow Export can take many multiples of Idle Timeout to empty the Cache after the measurement. This behavior is completely implementation-specific.

If the Collector has the capability to redirect the Flow Export data after the measurement time interval into a different capture buffer (or time stamp the received Flow Export data after that), this can be done in one step. Otherwise, each Flow Monitoring Throughput measurement at a certain packet rate needs to be executed twice -- once to capture the Flow Export data just for the measurement time

interval (to determine the actual Flow Export Rate) and a second time to capture all Flow Export data in order to determine Flow Record losses at that packet rate.

At the end of the measurement time interval, the DUT might still be processing Cache entries that belong to the Flows expired from the Cache before the end of the interval. These Flow Records might appear in an export packet sent only after the end of the measurement interval. This imprecision can be mitigated by use of large amounts of Flow Records during the measurement (so that the few Flow Records in one export packet can be ignored) or by use of timestamps exported with the Flow Records.

5.6. The Measurement Procedure

The measurement procedure is the same as the Throughput measurement in Section 26.1 of [RFC2544] for the traffic sending side. The DUT output analysis is done on the traffic generator receiving side for the test traffic, the same way as for RFC 2544 measurements.

An additional analysis is performed using data captured by the Collector. The purpose of this analysis is to establish the value of the Flow Export Rate during the current measurement step and to verify that no Flow Records were dropped during the measurement. The procedure for measuring the Flow Export Rate is described in Section 5.5.

The Flow Export performance can be significantly affected by the way the Flow monitoring implementation formats the Flow Records into the Flow Export packets. The ordering and frequency in which Control Information is exported and the number of Flow Records in one Flow Export packet are of interest. In the worst case scenario, there is just one Flow Record in every Flow Export packet.

Flow Export data should be sanity checked during the benchmark measurement for:

- a. the number of Flow Records per packet, by simply calculating the ratio of exported Flow Records to the number of Flow Export packets captured during the measurement (which should be available as a counter on the Collector capture buffer).
- b. the number of Flow Records corresponding to the export of Control Information per Flow Export packet (calculated as the ratio of the total number of such Flow Records in the Flow Export data and the number of Flow Export packets).

6. RFC 2544 Measurements

RFC 2544 measurements can be performed under two Flow monitoring setups (see also Section 3.4.2). This section details both and specifies ways to construct the test traffic so that RFC 2544 measurements can be performed in a controlled environment from the Flow monitoring point of view. A controlled Flow monitoring environment means that the tester always knows what Flow monitoring activity (Flow Export Rate) the traffic offered to the DUT causes.

This section is applicable mainly for the Throughput (RFC 2544, Section 26.1) and latency (RFC 2544, Section 26.2) measurements. It could also be used to measure frame loss rate (RFC 2544, Section 26.3) and back-to-back frames (RFC 2544, Section 26.4). Flow Export requires DUT resources to be generated and transmitted; therefore, the Throughput in most cases will be much lower when Flow monitoring is enabled on the DUT than when it is not.

Objective:

Provide RFC 2544 network device characteristics in the presence of Flow monitoring on the DUT. RFC 2544 studies numerous characteristics of network devices. The DUT forwarding and time characteristics without Flow monitoring present on the DUT can vary significantly when Flow monitoring is deployed on the network device.

Metric definition:

Metric as specified in [RFC2544].

The measured Throughput MUST NOT include the packet rate corresponding to the Flow Export data, because it is not user traffic forwarded by the DUT. It is generated by the DUT as a result of enabling Flow monitoring and does not contribute to the test traffic that the DUT can handle. Flow Export requires DUT resources to be generated and transmitted; therefore, the Throughput in most cases will be much lower when Flow monitoring is enabled on the DUT than when it is not.

6.1. Flow Monitoring Configuration

Flow monitoring configuration (as detailed in Section 4.3) needs to be applied the same way as discussed in Section 5 with the exception of the Active Timeout configuration.

The Active Timeout SHOULD be configured to exceed several times the measurement time interval (see Section 5.4). This ensures that if

measurements with two traffic components are performed (see Section 6.3.2), there is no Flow monitoring activity related to the second traffic component.

The Flow monitoring configuration does not change in any other way for the measurement performed in this section. What changes and makes the difference is the traffic configurations as specified in the sections below.

6.2. Measurements with the Flow Monitoring Throughput Setup

To perform a measurement with Flow Monitoring Throughput setup, the major requirement is that the traffic and Flow monitoring be configured in such a way that each sent packet creates one entry in the DUT Cache. This restricts the possible setups only to the measurement with two traffic components as specified in Section 6.3.2.

6.3. Measurements with a Fixed Flow Export Rate

This section covers the measurements where the RFC 2544 metrics need to be measured with Flow monitoring enabled, but at a certain Flow Export Rate that is lower than the Flow Monitoring Throughput.

The tester here has both options as specified in Sections 6.3.1 and 6.3.2.

6.3.1. Measurements with a Single Traffic Component

Section 12 of [RFC2544] discusses the use of protocol source and destination addresses for defined measurements. To perform all the RFC 2544 type measurements with Flow monitoring enabled, the defined Flow Keys SHOULD contain an IP source and destination address. The RFC 2544 type measurements with Flow monitoring enabled then can be executed under these additional conditions:

- a. the test traffic is not limited to a single, unique pair of source and destination addresses.
- b. the traffic generator defines test traffic as follows: it allows for a parameter to send N (where N is an integer number starting at 1 and is incremented in small steps) packets with source IP address A and destination IP address B before changing both IP addresses to the next value.

This test traffic definition allows execution of the Flow monitoring measurements with a fixed Flow Export Rate while measuring the DUT RFC 2544 characteristics. This setup is the better option since it

best simulates the live network traffic scenario with Flows containing more than just one packet.

The initial packet rate at N equal to 1 defines the Flow Export Rate for the whole measurement procedure. Subsequent increases of N will not change the Flow Export Rate as the time and Cache characteristics of the test traffic stay the same. This setup is suitable for measurements with Flow Export Rates below the Flow Monitoring Throughput.

6.3.2 Measurements with Two Traffic Components

The test traffic setup described in Section 6.3.1 might be difficult to achieve with commercial traffic generators or if the granularity of the traffic rates as defined by the initial packet rate at N equal to 1 are unsuitable for the required measurement. An alternative mechanism is to define two traffic components in the test traffic: one to populate Flow monitoring Cache and the second to execute the RFC 2544 measurements.

- a. Flow monitoring test traffic component -- the exact traffic definition as specified in Section 5.2.
- b. RFC 2544 Test Traffic Component -- test traffic as specified by RFC 2544 MUST create just one entry in the DUT Cache. In the particular setup discussed here, this would mean a traffic stream with just one pair of unique source and destination IP addresses (but could be avoided if Flow Keys were, for example, UDP/TCP source and destination ports and Flow Keys did not contain the addresses).

The Flow monitoring traffic component will exercise the DUT in terms of Flow activity, while the second traffic component will measure the RFC 2544 characteristics.

The measured Throughput is the sum of the packet rates of both traffic components. The definition of other RFC 1242 metrics remains unchanged.

7. Flow Monitoring Accuracy

The pure Flow Monitoring Throughput measurement described in Section 5 provides the capability to verify the Flow monitoring accuracy in terms of the exported Flow Record data. Since every Cache entry created in the Cache is populated by just one packet, the full set of captured data on the Collector can be parsed (e.g., providing the values of all Flow Keys and other Flow Record fields, not only the overall Flow Record count in the exported data), and each set of

parameters from each Flow Record can be checked against the parameters as configured on the traffic generator and set in packets sent to the DUT. The exported Flow Record is considered accurate if:

- a. all the Flow Record fields are present in each exported Flow Record.
- b. all the Flow Record fields' values match the value ranges set by the traffic generator (for example, an IP address falls within the range of the IP address increments on the traffic generator).
- c. all the possible Flow Record field values as defined at the traffic generator have been found in the captured export data on the Collector. This check needs to be offset against detected packet losses at the DUT during the measurement.

For a DUT with packet forwarding, the Flow monitoring accuracy also involves data checks on the received traffic, as already discussed in Section 4.

8. Evaluating Flow Monitoring Applicability

The measurement results, as discussed in this document and obtained for certain DUTs, allow for a preliminary analysis of a Flow monitoring deployment based on the traffic analysis data from the providers' network. An example of such traffic analysis in the Internet is provided by [CAIDA]; the way it can be used is discussed below. The data needed to estimate if a certain network device can manage the particular amount of live traffic with Flow monitoring enabled is:

Average packet size: 350 bytes
Number of packets per IP flow: 20

Expected data rate on the network device: 1 Gbit/s

The average number of Flows created per second in the network device is needed and is determined as follows:

$$\text{Flows per second} = \frac{\text{Expected packet rate}}{\text{Packet per flow}}$$

When using the above example values, the network device is required to process 18000 Flows per second. By executing the benchmarking as specified in this document, a platform capable of this processing can be determined for the deployment in that particular part of the user network.

Keep in mind that the above is a very rough and averaged Flow activity estimate, which cannot account for traffic anomalies; for example, a large number of DNS request packets that are typically small packets coming from many different sources and represent mostly just one packet per Flow.

9. Acknowledgements

This work was performed thanks to the patience and support of Cisco Systems NetFlow development team, namely Paul Aitken, Paul Atkins, and Andrew Johnson. Thanks to Benoit Claise for numerous detailed reviews and presentations of the document, and to Aamer Akhter for initiating this work. A special acknowledgment to the entire BMWG working group, especially to the chair, Al Morton, for the support and work on this document and Paul Aitken for a very detailed technical review.

10. Security Considerations

Documents of this type do not directly affect the security of the Internet or corporate networks as long as benchmarking is not performed on devices or systems connected to operating networks.

Benchmarking activities, as described in this memo, are limited to technology characterization using controlled stimuli in a laboratory environment, with dedicated address space and the constraints specified in sections above.

The benchmarking network topology will be an independent test setup and MUST NOT be connected to devices that may forward the test traffic into a production network, or misroute traffic to the test management network.

Further, benchmarking is performed on a "black-box" basis, relying solely on measurements observable external to the DUT.

Special capabilities SHOULD NOT exist in the DUT specifically for benchmarking purposes. Any implications for network security arising from the DUT SHOULD be identical in the lab and in production networks.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.

11.2. Informative References

- [RFC1242] Bradner, S., "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991.
- [RFC2285] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, February 1998.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC3917] Quittek, J., Zseby, T., Claise, B., and S. Zander, "Requirements for IP Flow Information Export (IPFIX)", RFC 3917, October 2004.
- [RFC3954] Claise, B., Ed., "Cisco Systems NetFlow Services Export Version 9", RFC 3954, October 2004.
- [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.
- [RFC5180] Popoviciu, C., Hamza, A., Van de Velde, G., and D. Dugatkin, "IPv6 Benchmarking Methodology for Network Interconnect Devices", RFC 5180, May 2008.
- [RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek, "Architecture for IP Flow Information Export", RFC 5470, March 2009.
- [RFC5695] Akhter, A., Asati, R., and C. Pignataro, "MPLS Forwarding Benchmarking Methodology for IP Flows", RFC 5695, November 2009.
- [CAIDA] Claffy, K., "The nature of the beast: recent traffic measurements from an Internet backbone", <http://www.caida.org/publications/papers/1998/Inet98/Inet98.html>

[IPFIX-CONFIG]

Muenz, G., Muenchen, TU, Claise, B., and P. Aitken,
"Configuration Data Model for IPFIX and PSAMP", Work in
Progress, July 2011.

[PSAMP-MIB] Dietz, T., Claise, B., and J. Quittek, "Definitions of
Managed Objects for Packet Sampling", Work in Progress,
October 2011.

[IPFIX-MIB] Dietz, T., Kobayashi, A., Claise, B., and G. Muenz,
"Definitions of Managed Objects for IP Flow Information
Export", Work in Progress, March 2012.

Appendix A. (Informative) Recommended Report Format

Parameter	Units
-----	-----
Test Case	test case name (Sections 5 and 6)
Test Topology	Figure 2, other
Traffic Type	IPv4, IPv6, MPLS, other
Test Results	
Flow Monitoring Throughput	Flow Records per second or Not Applicable
Flow Export Rate	Flow Records per second or Not Applicable
Control Information Export Rate Throughput (Other RFC 1242 Metrics)	Flow Records per second packets per second (as appropriate)
General Parameters	
DUT Interface Type	Ethernet, POS, ATM, other
DUT Interface Bandwidth	MegaBits per second
Traffic Specifications	
Number of Traffic Components	(see Sections 6.3.1 and 6.3.2)
For each traffic component:	
Packet Size	bytes
Traffic Packet Rate	packets per second
Traffic Bit Rate	MegaBits per second
Number of Packets Sent	number of entries
Incremented Packet Header Fields	list of fields
Number of Unique Header Values	number of entries
Number of Packets per Flow	number of entries
Traffic Generation	linearly incremented or randomized
Flow monitoring Specifications	
Direction	ingress, egress, both
Observation Points	DUT interface names
Cache Size	number of entries
Active Timeout	seconds
Idle Timeout	seconds
Flow Keys	list of fields
Flow Record Fields	total number of fields
Number of Flows Created	number of entries
Flow Export Transport Protocol	UDP, TCP, SCTP, other
Flow Export Protocol	IPFIX, NetFlow, other
Flow Export data packet size	bytes
Flow Export MTU	bytes

Parameter	Units (continued)
MPLS Specifications	(for traffic type MPLS only)
Tested Label Operation	imposition, swap, disposition

The format of the report as documented in this appendix is informative, but the entries in the contents of it are required as specified in the corresponding sections of this document.

Many of the configuration parameters required by the measurement report can be retrieved from the [IPFIX-MIB] and [PSAMP-MIB] MIB modules, and from the [IPFIX-CONFIG] YANG module or other general MIBs. Therefore, querying those modules from the DUT would be beneficial: first of all, to help in populating the required entries of the measurement report, and also to document all the other configuration parameters from the DUT.

Appendix B. (Informative) Miscellaneous Tests

This section lists tests that could be useful to assess a proper Flow monitoring operation under various operational or stress conditions. These tests are not deemed suitable for any benchmarking for various reasons.

B.1. DUT Under Traffic Load

The Flow Monitoring Throughput should be measured under different levels of static traffic load through the DUT. This can be achieved only by using two traffic components as discussed in Section 6.3.2. One traffic component exercises the Flow Monitoring Plane. The second traffic component loads only the Forwarding Plane without affecting Flow monitoring (i.e., it creates just a certain amount of permanent Cache entries).

The variance in Flow Monitoring Throughput as a function of the traffic load should be noted for comparison purposes between two DUTs of similar architecture and capability.

B.2. In-Band Flow Export

The test topology in Section 4.1 mandates the use of a separate Flow Export interface to avoid the Flow Export data generated by the DUT to mix with the test traffic from the traffic generator. This is necessary in order to create clear and reproducible test conditions for the benchmark measurement.

The real network deployment of Flow monitoring might not allow for such a luxury -- for example, on a very geographically large network.

In such a case, the Flow Export will use an ordinary traffic forwarding interface, e.g., in-band Flow Export.

The Flow monitoring operation should be verified with in-band Flow Export configuration while following these test steps:

- a. Perform the benchmark test as specified in Section 5. One of the results will be how much bandwidth Flow Export used on the dedicated Flow Export interface.
- b. Change Flow Export configuration to use the test interface.
- c. Repeat the benchmark test while the receiver filters out the Flow Export data from analysis.

The expected result is that the Throughput achieved in step a. is same as the Throughput achieved in step c. provided that the bandwidth of the output DUT interface is not the bottleneck (in other words, it must have enough capacity to forward both test and Flow Export traffic).

B.3. Variable Packet Size

The Flow monitoring measurements specified in this document would be interesting to repeat with variable packet sizes within one particular test (e.g., test traffic containing mixed packet sizes). The packet forwarding tests specified mainly in [RFC2544] do not recommend performing such tests. Flow monitoring is not dependent on packet sizes, so such a test could be performed during the Flow Monitoring Throughput measurement, and verification of its value does not depend on the offered traffic packet sizes. The tests must be carefully designed in order to avoid measurement errors due to the physical bandwidth limitations and changes of the base forwarding performance with packet size.

B.4. Bursty Traffic

RFC 2544, Section 21 discusses and defines the use of bursty traffic. It can be used for Flow monitoring testing to gauge some short-term overload DUT capabilities in terms of Flow monitoring. The test benchmark here would not be the Flow Export Rate the DUT can sustain, but the absolute number of Flow Records the DUT can process without dropping any single Flow Record. The traffic setup to be used for this test is as follows:

- a. each sent packet creates a new Cache entry.
- b. the packet rate is set to the maximum transmission speed of the DUT interface used for the test.

B.5. Various Flow Monitoring Configurations

This section translates the terminology used in the IPFIX documents ([RFC5470], [RFC5101], and others) into the terminology used in this document. Section B.5.2 proposes another measurement that is impossible to verify in a black box test manner.

B.5.1. Throughput without the Metering Process

If the Metering Process is not defined on the DUT it means no Flow monitoring Cache exists and no Flow analysis occurs. The performance measurement of the DUT in such a case is just pure [RFC2544] measurement.

B.5.2. Throughput with the Metering Process

If only the Metering Process is enabled, Flow analysis on the DUT is enabled and operational but no Flow Export happens. The performance measurement of a DUT in such a configuration represents a useful test of the DUT's capabilities (this corresponds to the case when the network operator uses Flow monitoring, for example, for manual detection of denial-of-service attacks, and does not wish to use Flow Export).

The performance testing on this DUT can be performed as discussed in this document, but it is not possible to verify the operation and results without interrogating the DUT.

B.5.3. Throughput with the Metering and Exporting Processes

This test represents the performance testing as discussed in Section 6.

B.6. Tests With Bidirectional Traffic

Bidirectional traffic is not part of the normative benchmarking tests based on discussion with and recommendation of the Benchmarking working group. The experienced participants stated that this kind of traffic did not provide reproducible results.

The test topology in Figure 2 can be expanded to verify Flow monitoring functionality with bidirectional traffic using the interfaces in full duplex mode, e.g., sending and receiving simultaneously on each of them.

The same rules should be applied for Flow creation in the DUT Cache (as per Sections 4.1 and 4.3.1) -- traffic passing through each Observation Point should always create a new Cache entry in the

Cache, e.g., the same traffic should not be just looped back on the receiving interfaces to create the bidirectional traffic flow.

B.7. Instantaneous Flow Export Rate

Additional useful information when analyzing the Flow Export data is the time distribution of the instantaneous Flow Export Rate. It can be derived during the measurements in two ways:

- a. The Collector might provide the capability to decode Flow Export during capturing and at the same time count the Flow Records and provide the instantaneous (or simply, an average over shorter time interval than specified in Section 5.4) Flow Export Rate.
- b. The Flow Export protocol (like IPFIX [RFC5101]) can provide time stamps in the Flow Export packets that would allow time-based analysis and calculate the Flow Export Rate as an average over much shorter time interval than specified in Section 5.4.

The accuracy and shortest time average will always be limited by the precision of the time stamps (1 second for IPFIX) or by the capabilities of the DUT and the Collector.

Author's Address

Jan Novak (editor)
Cisco Systems
Edinburgh
United Kingdom
EMail: janovak@cisco.com

