IP Flow Information Accounting and Export Benchmarking Methodology
draft-ietf-bmwg-ipflow-meth-01.txt

Abstract

   This document provides methodology and framework for quantifying
   performance impact of monitoring of IP flows on a network device and
   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 the
   Architecture for IP Flow Information Export [RFC5470].

Status of this Memo

Copyright Notice

the Trust Legal Provisions and are provided without warranty as
described in the Simplified BSD License.

Conventions used in this document

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

Table of Contents

1. Introduction

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

   Flow monitoring is defined in the Architecture for IP Flow
   Information Export [RFC5470] and related IPFIX documents.

   What is the cost of enabling the IP Flow monitoring and export to a
   collector ? This is the basic question that this methodology is
   designed to answer.

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

   Since Flow monitoring will in most cases run on network devices also
   forwarding packets, the methodology for RFC2544 measurements (with
   IPv6 and MPLS specifics defined in [RFC5180] and [RFC5695]
   respectively) in the presence of Flow monitoring is also employed
   here.

   The most significant performance parameter is the rate at which IP
   flows are created and expired in the network devices memory and
   exported to a collector. Therefore, this document focuses on a
   methodology on how to measure the maximum IP flow rate that a
   network device can sustain without impacting the forwarding plane,
   without losing any IP flow information, and without compromising the
   IP flow accuracy.

   [RFC2544], [RFC5180] and [RFC5695] specify benchmarking of network
   devices forwarding IPv4, IPv6 and MPLS [RFC3031] traffic,
   respectively. The methodology specified here stays the same for any
   traffic type. The only restriction is the actual Flow monitoring

support for the particular traffic type.

A variety of different network device architectures exist that are
capable of Flow monitoring and export. As such, this document does

not attempt to list the various white box variables (CPU load,
memory utilization, TCAM utilization etc) that could be gathered as
they do 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 mostly based on [RFC5470],
[RFC2285] and [RFC1242] as summarised in the section 2.1. The only
new terms needed for this methodology are defined in the following
section 2.2.

2.1 Existing Terminology

Device Under Test (DUT)   [RFC2285, section 3.1.1]

Flow   [RFC5470, section 2]

Flow Key   [RFC5470, section 2]

Flow Record   [RFC5470, 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]

Packet Sampling   [RFC5476, section 2]

2.2 New Terminology

2.2.1 Cache

Definition:

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

2.2.2 **Cache Size**

   Definition:
      The size of the Cache in terms of how many entries of Flow
      Records 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 Flow Records

2.2.3 **Active Timeout**

   Definition:
      For long-running Flows, the time interval after which the Metering
      Process expires a Flow Record from the Cache so that only regular
      Flow updates are exported.

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

      Flows are considered long-running when they last longer than
      several multiples of the Active Timeout or contain larger amount
      of packets (in the case of Active Timeout is zero) than usual for
      a single transaction based Flows, in the order of tens of packets
      and higher.

   Measurement units:
      Seconds

2.2.4 **Inactive Timeout**

   Definition:
      The time interval used by the Metering Process to expire a Flow
      Record from the Cache, when no more packets belonging to that
      specific Flow are observed during the interval.

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

   Measurement units:
      Seconds

2.2.5 **Flow Export Rate**

   Definition:
      Number of Flow Records that expire from the Cache (as defined by
      the Flow Expiration term) and are exported to the Collector within
      a measurement time interval.

      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 nearly instantaneous export
      rate rather than average export rate on the DUT) is limited by the
      export capabilities of the particular Flow monitoring
      implementation.

   Measurement units:
      Number of Flow Records per second

3. **Flow Monitoring Performance Metric**

3.1 **The Definition**

   Flow Monitoring Throughput

   Definition:
      The maximum Flow Export Rate the DUT can sustain without losing a
      single Flow Record expired from the Cache. Additionally, for the
      packet forwarding devices, also 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

3.2 **Device Applicability**

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

The Collector performance is out of scope of this document.

3.3 **Measurement Concept**

   The traffic in the Figure 1 represents the test traffic sent to the
   DUT and forwarded by the DUT. When testing devices which do not act
   as network packet forwarding devices (appliances - probes, sniffers,
   replicators) the forwarding plane is simply an Observation Point as
   defined in section 2 of [RFC5470]. The RFC2544 Throughput of such
   devices will simply be always zero and the only applicable
   performance metrics is Flow Monitoring Throughput.

   The Flow monitoring enabled (see section 4.3) on the DUT (and
   represented in the Figure 1 by the Flow Monitoring Plane) uses the
   traffic information provided by the Forwarding Plane and configured
   Flow Keys to create the Flow Records representing the traffic
   forwarded (or observed) by the DUT. The Flow Records are stored in
   the Flow monitoring Cache and expired from there depending on the
   Cache configuration (Active and Inactive Timeouts, number of Flow
   Records and the Cache Size) and the traffic pattern. The expired Flow
   Records are exported from the DUT to the Collector (see Figure 2 in
   section 4).

```
                 +------------------------ +
                 | IPFIX | Netflow | Others |
                 +------------------------ +
                 |            ^           |
                 |            ^           |
                 |       Flow Export      |
                 |            ^           |
                 |            ^           |
                 |     +-------------+     |
                 |     |    Flow     |     |
                 |     | Monitoring  |     |
                 |     |   Plane     |     |
                 |     +-------------+     |
                 |            ^           |
                 |            ^           |
                 |     traffic information |
                 |            ^           |
                 |            ^           |
                 |     +-------------+     |
                 |     |             |     |
    traffic ---|---->| Forwarding  |------|---->
                 |     |   Plane     |     |
                 |     +-------------+     |
                 |                        |
                 |           DUT          |
                 +------------------------ +
```
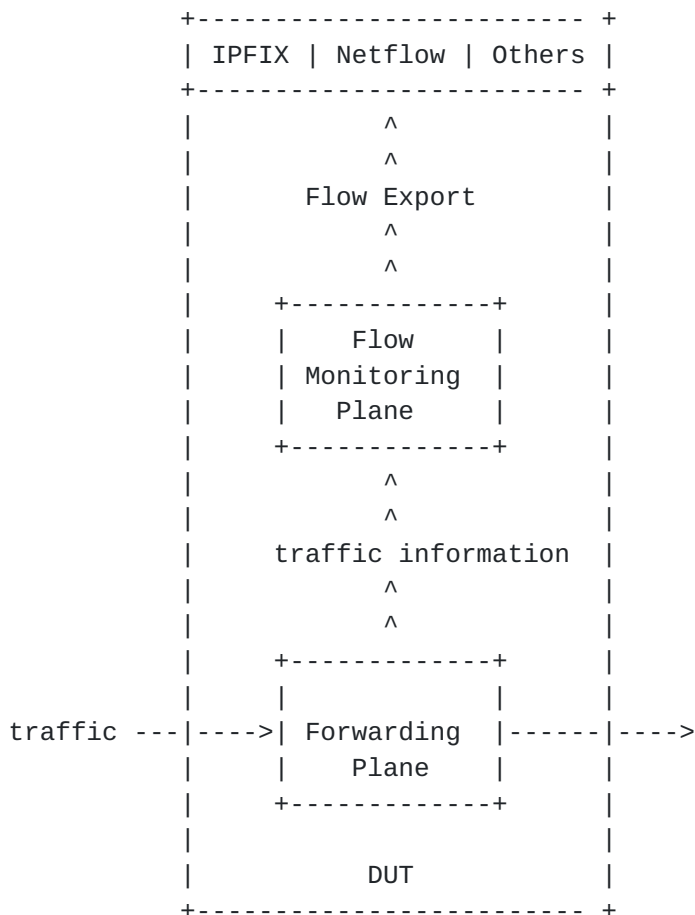
Figure 1. The functional block diagram of the DUT

The Forwarding Plane and Flow Monitoring Plane represent two separate
functional blocks, each with it's own performance capability. The
Forwarding Plane handles user data packets and is fully characterised
by the metrics defined by [RFC2544].

The Flow Monitoring Plane handles Flow Records which reflect the forwarded traffic. The metric that measures the Flow Monitoring Plane performance is 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 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 Device Under Test (DUT) with Flow monitoring enabled contains two functional blocks which 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 Flow Monitoring Plane and Forwarding Plane (see Figure 1) need to be looked at as two independent blocks (and the performance of each of them measured independently) but on the other hand when measuring the performance of one of them the status and performance of the other MUST be known and benchmarked when both are present.

3.4.1 Flow Monitoring Plane Performance Measurement

The Flow Monitoring Throughput MUST be (and can only be) measured with one packet per Flow as specified in the section 5. This traffic type represents the most demanding traffic from the Flow monitoring point of view and will exercise the Flow Monitoring Plane (see Figure 1) of the DUT most. The exit criteria for the Flow Monitoring Throughput measurement are one of the following (e.g. if any of the conditions is reached):

a. The Flow Export Rate at which the DUT starts to drop Flow Records 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)

3.4.2 Forwarding Plane Performance Measurement

The Forwarding Plane (see Figure 1) performance metrics are fully specified by [RFC2544] and MUST be measured accordingly. A detailed traffic analysis (see below) with relation to Flow monitoring MUST be performed prior of any RFC2544 measurements. Mainly the Flow Export

Rate caused by the test traffic during an [RFC2544](#) measurement MUST
   be known and reported.

The required 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 RFC2544 performance metrics can be measured in one of the three
modes:

a. As a baseline of forwarding performance without Flow monitoring
b. At certain level of Flow monitoring activity specified by a Flow
   Expiration Rate lower than Flow Monitoring Throughput
c. At the maximum 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 RFC2544. The details how
to setup the measurements in points b. and c. are in the section 6.

4. Measurement Set Up

This section concentrates on the set-up 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. Generally, those appliances instead of the Forwarding Plane
will have some kind of feed (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.

The measurement set-up is identical to the one 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 generated from the sender to the receiver, where the received
traffic is analyzed to check it is identical to the generated
traffic.

The ideal way to implement the measurement is using one traffic
generator (device providing the sender and receiver capabilities)
with a sending port and a receiving port. This allows for an easy
check if all the traffic sent by the sender was transmitted by the
DUT and received at the receiver.

```
                        +-----------+
                        |           |
                        | Collector |
                        |           |
                        |Flow Record|
                        | analysis  |
                        |           |
                        +-----------+
                             ^
                             | Flow Export
                             |
                             | Export Interface
        +--------+       +-------------+          +----------+
        |        |       |             |          |          |
        |        |    (*)|             |          | receiver |
        | sender |------>|     DUT     |--------->|          |
        |        |       |             |          | traffic  |
        |        |       |             |          | analysis |
        +--------+       +-------------+          +----------+
```
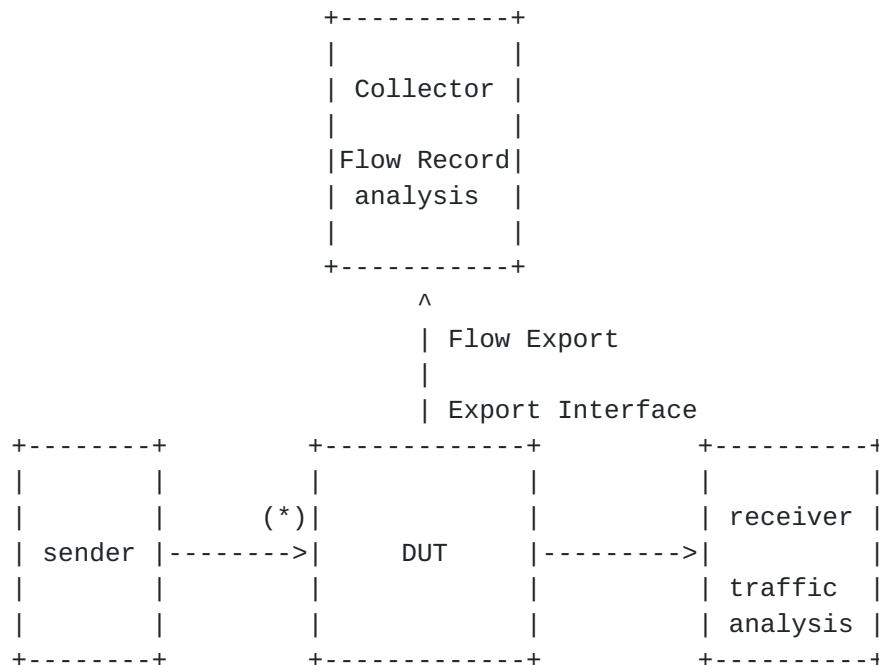
Figure 2 Measurement topology with unidirectional traffic

The export interface (connecting the Collector) MUST NOT be used for
forwarding the 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.

Note that more complex topologies might be required. For example, if
the effects of enabling Flow monitoring on several interfaces are of
concern or the media maximum speed 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 Set Up**

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

The base 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
measurements results to take into account the possible change of
packets 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) MUST be configured with
sufficient output buffers to avoid dropping the Flow Export data due
to a simple lack of resources in the interface hardware.

[4.3](#) Flow Monitoring Configuration

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

The DUT MUST support Flow monitoring architecture as specified by
[RFC5470]. The DUT SHOULD support IPFIX [RFC5101] for easier results
comparison.

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

4.3.1 Observation Points

The Observation Points specify the interfaces and direction where
the Flow monitoring traffic analysis is 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(s) only
   b. egress port(s) /ports only
   c. both ingress and egress

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

In the case when both ingress and egress Flow monitoring is
enabled on one DUT the results analysis needs to take into account
that each Flow will be represented in the DUT Cache by two Flow
Records (one for each direction) and therefore also the Flow
Export will 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 which do not overlap, e.g. each packet (or set of packets
if measuring with more than one packet per Flow) 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 report of the results.

4.3.2 Metering Process

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

   Cache Size available to the DUT operation MUST be known and taken
   into account when designing the measurement as specified in the
   section 5.

   Inactive and Active Timeouts MUST be known and taken into account
   when designing the measurement as specified in the section 5.

   The Cache Size, the Inactive and Active Timeouts, and if present,
   the specific Packet Sampling techniques and associated parameters
   MUST be included as part of the results report.

4.3.3 Exporting Process

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

   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 composed of
   the Inactive and Active Timeouts and Cache Size.

   The Exporting Process SHOULD be configured with IPFIX [RFC5101] as
   the protocol to use to format the Flow Export data. If the Flow
   monitoring implementation does not support it, proprietary
   protocols MAY be used.

   Various Flow monitoring implementations might use different
   default values regarding the export of Control Information. The
   Flow Export corresponding to Control Information SHOULD be
   analyzed and reported as a separate item on the measurement
   report. Preferably, the export of Control Information SHOULD
   always be configured same.

   IPFIX documents [RFC5101] in section 10 and [RFC5470] in section
   8.1 discuss the possibility to deploy 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 same transport layer
   protocol SHOULD be compared. If the Flow monitoring implementation
   allows to use all of UDP, TCP and SCTP as the transport layer
   protocols, each of the protocols SHOULD be measured in a separate
   measurement run.

   If reliable transport protocol is used for the transmission of the

Flow Export data from DUT, the configuration of the transport
session MUST allow for non-blocking data transmission. An example
of parameters to look at would be TCP window size or maximum
segment size (MSS).

4.3.4 Flow Records

Flow Record defines the traffic parameters which Flow monitoring
uses to analyze the traffic and MUST be configured in order to
perform the analysis. The Flow Key fields of the Flow Record
define the traffic parameters which will be used to create new
Flow Records in the DUT Cache.

The Flow Record definition is implementation specific. A Flow
monitoring implementation might allow for only fixed Flow Record
definition, based on the most common IP parameters in the IPv4 or
IPv6 headers - like source and destination IP addresses, IP
protocol numbers or transport level port numbers. Another
implementation might allow the user to actually define his own
completely arbitrary Flow Record to monitor the traffic. The
requirement for the measurements defined in this document is only
the need for a large number of Flow Records in the Cache. The Flow
Keys needed to achieve that will typically be source and
destinations IP addresses and transport level port numbers.

Recommended full IPv4, IPv6 or MPLS Flow Record:
   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

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

   Flow Keys
           Source IP address
           Destination IP address

   Other fields
           Packet counter

or:

   Flow Key fields
           Transport layer source port
           Transport layer destination port

Other fields
              Packet counter

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 can
and MUST be compared only for exactly same Flow Record
configuration.

4.3.5 MPLS Measurement Specifics

The Flow Record configuration for measurements with MPLS
encapsulated traffic SHOULD contain MPLS label or any other field
which is part of the MPLS header.

The DUT Cache SHOULD be checked prior the performance measurement to
contain the correct MPLS related information.

The captured export data at the Collector SHOULD be checked for the
presence of MPLS labels or the monitored MPLS parameters. MPLS
forwarding performance document [RFC5695] specifies number of
possible MPLS label operations to test. The Observation Points
SHOULD 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 SHOULD be configured in such a way, that all the traffic is
subject of the measured MPLS label operation.

4.4 Collector

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

The Collector can be used as exclusively 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.

However if the Collector is also used to decode the Flow Export data
then it SHOULD support IPFIX [RFC5101] for easier 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 to capture at the full rate the export
packets sent from the DUT without losing any of them. In the case of
the use of 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 Collector SHOULD support Ethernet type of interface to connect to
the DUT but any media which allows data capturing and analysis can be
used.

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

4.5 Packet Sampling

A Flow monitoring implementation might provide the capability to
analyze the Flows after Packet Sampling is performed. The possible
procedures and ways of Packet Sampling are described in [RFC5476]
and [RFC5475] and only those SHOULD be used for measurements.

If the DUT is configured with one of the sampling techniques as
specified in [RFC5475] the measurement report MUST include this
sampling technique along with its parameters. The presence of the
configured sampling technique on the DUT and its parameters SHOULD be
verified in the Flow Export data as received on the Collector.

Packet Sampling will affect the measured Flow Export Rate. If
systematic sampling (see section 6.5 of [RFC5476]) is in use, the
Flow Export Rate can be derived from the packet rates (see section 5
of this document) using the configured sampling parameters. If random
sampling is in use the Flow Export Rate can be derived from the
traffic rates as obtained on the receiver side of the traffic
generator, provided that packet losses can be excluded by monitoring
the DUT forwarding statistics.

If measurements are performed with Flows containing more than one
packet per Flow (see section 6.4 of this document) the sampling ratio
SHOULD always be higher than the number of packets in the Flows (for
small number of packets per Flow). This significantly decreases the
probability of erasing a whole Flow to a minimum and the measured
Flow Expiration Rate stays unaffected by sampling.

If Flow accuracy analysis (see section 7) is performed, the results
will be always affected by Packet Sampling and the complete check of
data cannot be performed.

This document does not intend to study the effects of Packet Sampling
itself on the network devices but Packet Sampling can simply be
applied as part of the Flow monitoring configuration on the DUT and
perform the measurements as specified in the later sections.
Consideration needs to be made when evaluating measurements results
to take into account the change of packet rates offered to the DUT
and especially to Flow monitoring after Packet Sampling is applied.

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.

The most common transmission media and corresponding frame formats
(Ethernet, Packet over SONET) for IPv4, IPv6 and MPLS traffic are
specified within [RFC2544], [RFC5180] and [RFC5695].

4.7 Frame Sizes

Frame sizes of the traffic analyzed by the to use are specified in [RFC2544] section 9 for Ethernet type interfaces (64, 128, 256, 1024, 1280, 1518 bytes) and in [RFC5180] section 5 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 which could negatively affect measured performance values.

4.8 Flow Export Data Packet Sizes

The Flow monitoring performance will be affected by the packet size the particular implementation uses to transmit Flow Export data to the Collector. The used packet size SHOULD be part of the test 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 media largest available value.

4.9 Illustrative Test Set-up Examples

The below examples represent only hypothetical test set-up to clarify the use of Flow monitoring parameters and configuration together with traffic parameters to test Flow monitoring. The actual benchmarking specifications are in the sections 5 and 6.

4.9.1 Example 1 - Inactive Timeout Flow Expiration

The traffic generator sends 1000 packets per second in 10000 defined streams, each stream identified by an unique destination IP address. Each stream has then packet rate 0.1 packets per second. The packets are sent in a round robin fashion (stream 1 to 10000) while incrementing the destination IP address with each sent packet.

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

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

A packet with destination IP address equal to A is sent every 10 seconds, so it means that the Flow Record is refreshed in the Cache every 10 seconds, while the Inactive Timeout is 5 seconds. In this case the Flow Records will expire from the Cache due to the Inactive Timeout and when a new packet is sent with the same IP address A it will create a new Flow Record in the Cache.

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 Flow Record and we send 1000 packets per second.

The expected number of Flow Record entries in the Cache during the
whole measurement is around 5000. It corresponds to the Inactive
Timeout being 5 seconds and during those five seconds 5000 entries
are created. This expectation might change in real measurement
set-ups witch large Cache Sizes and high packet rates where the
export rate might be limited and lower than the offered Flow Export
Rate. This behaviour 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 an unique destination IP address.
Each stream has then packet rate 10 packets per second. The packets
are sent in a round robin fashion while incrementing (stream 1 to
100) the destination IP address with each sent packet.

The configured Cache Size is 1000 Flow Records. The configured
Active Timeout is 100 seconds, the Inactive Timeout is 10 seconds.

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

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

A packet with destination IP address equal to A is sent every 0.1
second, so it means that the Flow Record is refreshed in the Cache
every 0.1 second, while the Inactive Timeout is 10 seconds. In this
case the Flow Records will not expire from the Cache until the Active
Timeout, e.g. they will expire every 100 seconds and then the Flow
Records 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 no Flow Records 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 2/3 of Flow
Record per second since during the 290 seconds period we expired 2
times the same 100 of Flows.

5. Flow Monitoring Throughput Measurement Methodology

    Objective:

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

    Metric definition:

        Flow Monitoring Throughput - see section 3.

    Discussion:

        The Flow monitoring implementations might chose to handle
        differently Flow Export from a partially empty Cache or in the
        situation when the Cache is fully occupied by the Flow Records.
        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 abstract from all the
        possible behaviours and define one measurement procedure covering
        all the possibilities. 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 Flow Keys sets that the traffic
        generator (sender) provides should be multiple times larger than
        the Cache Size. This way the Flow Records in the Cache never get
        updated before Flow Expiration and Flow Export. The Cache Size
        MUST be known in order to define the measurements circumstances
        properly.

    Inactive Timeout
        Inactive Timeout is set (if configurable) to the minimum possible
        value on the network device. This makes sure the Flow Records are
        expired as soon as possible and exported out of the DUT Cache. It
        MUST be known in order to define the measurements circumstances
        completely and equally across implementations.

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

Flow Keys Definition:
    Needs to allow for large numbers of unique Flow Records to be
    created in the Cache 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
    Flow Record in the Cache.

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 Flow Record
    in the DUT Cache.

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

  Measurement Duration
    The measurement duration MUST be at least two times longer than
    the Inactive Timeout otherwise no Flow Export would be seen. The
    measurement duration SHOULD guarantee that the number of Flow
    Records created during the measurement exceeds the available Cache
    Size on the DUT.

5.3 Cache Population

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

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

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

5.4 Measurement Time Interval

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

RFC2544 specifies with the precision of the packet beginning and end
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
time start and stop 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) - (Inactive Timeout).

This supposes that the Cache Size is large enough so that the time to
fill it up with Flow Records is longer than Inactive Timeout.
Otherwise the time to fill up the Cache needs to be used for
calculation of the measurement time interval in the place of the
Inactive Timeout.

Instead of measuring the absolute values of stop and start time it is
possible to setup the traffic generator to send traffic for certain
pre-defined 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 Inactive Timeout (or the time needed to fill up the Cache) causes
delay of the Flow Export data behind the test traffic which is
forwarded by the DUT. E.g. if the traffic starts at time point X Flow
Export will start only at the time point X + Inactive Timeout (or X +
time to fill up the Cache). Since Flow Export capture needs to stop
with the traffic (because that's when the DUT stops to process the
Flow Records at the given rate) the time interval during which the
DUT kept exporting data is by Inactive Timeout shorter 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 Flow Record drops 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 Flow Records at the rate the packets are sent. When
      traffic generation finishes, the behaviour when emptying the Cache
      is completely implementation specific and the Flow Export data from
      this period cannot be therefore used for the benchmarking.
   b. In the second step all the Flow Export data from the DUT MUST be
      captured in order to be capable 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 and higher)
      are in use the Flow Export can take many multiples of Inactive
      Timeout to empty the Cache after the measurement. This behaviour is
      completely implementation specific.

   If the Collector has the capability to redirect the Flow Export data
   after the measurement time interval into 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
   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 Expiration Rate) and second time to capture all Flow
   Export data in order to determine Flow Record losses at that packet
   rate.

   This Flow Export Rate procedure is fully applicable to all
   measurement set-ups but can be simplified for the cases with high
   cache population (see section 5.3) when the Cache is filled up with
   Flow Records within first few seconds of the measurement. In such a
   case the DUT has no choice but to process all the Flows at the
   incoming packet rate and the Flow Export Rate is
   numerically equal to the packet rate. Thus only step b. really needs
   to be performed.

5.6 **The Measurement Procedure**

   The measurement procedure is same as the Throughput measurement in
   the 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 RFC2544 measurements.

An additional analysis is performed using data captured by the
Collector. The purpose of this analysis is to establish the value of
Flow Export Rate during the current measurement step and to verify

that no Flow Records were dropped during the measurement. The
procedure to measure Flow Export Rate is described in the 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 in terms of ordering and frequency of Control
Information export and mainly the number of Flow Records in one Flow
Export packet. The worst case scenario here 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 and 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 Control Information Flow Records 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). It should be several orders of magnitude less than one
   Flow Record per Flow Export packet or at most in some special
   configuration one unique set of Control Data in each Flow Export
   packet.

6. RFC2544 Measurements

RFC2544 measurements can be performed under two Flow Monitoring set-
ups (see also section 3.4.2). This section details both of them and
specifies the ways how to construct the test traffic so that RFC2544
measurements can be performed in a controlled environment also from
the Flow monitoring point of view. 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 RFC2544 throughput (RFC2544
section 26.1) and latency (RFC2544 section 26.2)measurement. It
could be used also to measure frame loss rate (RFC2544 section 26.3)
and back-to-back frames (RFC2544 section 26.4). It is irrelevant for
the rest of RFC2544 network interconnect devices characteristics.

Objective:

   Provide RFC2544 network device characteristics in the presence of
   Flow monitoring on the DUT. The RFC2544 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 starts to be deployed on

the network device.

Metric definition:

  Metric as specified in [RFC2544].

The measured RFC2544 Throughput MUST NOT include the packet rate
corresponding to the Flow Export data, because it is control type
traffic, generated by the DUT as a result of enabling Flow monitoring
and does not contribute to the test traffic which the DUT can handle.
It requires DUT resources to be generated and transmitted and
therefore the RFC2544 Throughput will be in most cases much lower
when Flow monitoring is enabled on the DUT than without it.

6.1 Flow Monitoring Configuration

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

The Active Timeout SHOULD be configured to exceed several times the
measurement time interval (see section 5.4). This makes sure that if
the measurements with two traffic components are performed (see
section 6.5) 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 Set-up

The major requirement to perform a measurement with Flow Monitoring
Throughput set-up is that the traffic and Flow monitoring is
configured in such a way that each sent packet creates one Flow
Record in the DUT Cache. This restricts the possible set-ups only to
the measurement with two traffic components as specified in the
section 6.5.

Note that for software based platforms (as already discussed in
Section 3.5) the two traffic components set-up might not be
necessary. This is to certain extent implementation specific. The two
traffic components set-up on software based platforms can still be
used to perform the type of measurements as discussed in the section
B.1.

6.3 Measurements With Fixed Flow Expiration Rate

This section covers the measurements where the RFC2544 metrics need

to be measured with Flow monitoring enabled but at certain Flow
Export Rate lower than Flow Monitoring Throughput.

The tester here has both options as specified in the [section 6.4](#) and
6.5.

6.4 Measurements With Single Traffic Component

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

a. the test traffic is not limited to single unique pair of source
   and destination address
b. the traffic generator defines test traffic as follows:
   allow for a parameter to say send N (where N is an integer number
   starting at 1 and 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 fixed Flow Export Rate while measuring the DUT
RFC2544 characteristics. This set-up 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 Expiration
Rate for the whole measurement procedure. The consequent increases
of N will not change Flow Expiration Rate as the time and Cache
characteristics of the test traffic stay the same. This set-up is
suitable for measurements with Flow Export Rates below the Flow
Monitoring Throughput.

6.5 Measurements With Two Traffic Components

The test traffic set-up in the section 6.4 might be difficult to
achieve with commercial traffic generators or the granularity of the
traffic rates as defined by the initial packet rate at N equal to 1
might not be suitable for the required measurement. An alternate
mechanism is to define two traffic components in the test traffic.
One to populate Flow monitoring Cache and the second one to execute
the RFC2544 measurements.

a. Flow monitoring test traffic component - the exact traffic
   definition as specified in the section 5.2.
b. RFC2544 Test Traffic Component - test traffic as specified by
   RFC2544 MUST create just one Flow Record in the DUT Cache. In
   the particular set-up 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
RFC2544 characteristics.

The measured RFC2544 Throughput is the sum of the packet rates of
both traffic components, the definition of other RFC2544 metrics
remains unchanged.

7. Flow Monitoring Accuracy

The pure Flow monitoring measurement in section 5 provides the
capability to verify the Flow monitoring accuracy in terms of the
exported Flow Record data. Since every Flow Record 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 packet 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 as set by
   the traffic generator (for example an IP address falls within the
   range of the IP addresses increments on the traffic generator)
c. all the possible Flow Record fields values as defined at the
   traffic generator have been found in the captured export data
   on the Collector. This check needs to be offset to potential
   detected packet losses at the DUT during the measurement

If Packet Sampling is deployed then only verifications in point a.
and b. above can be performed.

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] and the way it can be used is discussed below.
The data needed to make an 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

This results in:

Expected packet rate: 357 000 pps

being (1 Gbit/s divided by 350 bytes/packet)

Flows per second: 18 000

being (packet rate 357 000 pps divided by 20 packets per IP Flow)

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

9. Acknowledgements

This work could have been performed thanks to the patience and
support of Cisco Systems Netflow development team, namely Paul
Aitken, Paul Atkins and Andrew Johnson. Thanks belong to Benoit
Claise for numerous detailed reviews and presentations of the
document and Aamer Akhter for initiating this work.

10. IANA Considerations

This document makes no requests of IANA.

11. 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 the 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.

12. References

12.1. Normative References

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

   [RFC2544] Bradner, S., "Benchmarking Methodology for Network
              Interconnect Devices", Informational, RFC 2544, April 1999

   [RFC5470] Sadasivan, G., Brownlee, N., Claise, B., and J. Quittek,
              "Architecture Model for IP Flow Information Export",
              RFC 5470, April 2011

12.2. Informative References

   [RFC1242] Bradner, S., "Benchmarking Terminology for Network
              Interconnection Devices", RFC 1242, July 1991

   [RFC2285] Mandeville R., "Benchmarking Terminology for LAN Switching
              Devices", Informational, RFC 2285, November 1998

   [RFC3031] E. Rosen, A. Viswanathan, R. Callon, "Multiprotocol Label
              Switching Architecture", Standards Track, RFC 3031,
              January 2001

   [RFC3917] Quittek J., "Requirements for IP Flow Information Export
              (IPFIX)", Informational, RFC 3917, October 2004.

   [RFC5101] Claise B., "Specification of the IP Flow Information
              Export (IPFIX) Protocol for the Exchange of IP Traffic
              Flow Information", Standards Track, RFC 5101, January 2008

   [RFC5180] C. Popoviciu, A. Hamza, D. Dugatkin, G. Van de Velde,
              "IPv6 Benchmarking Methodology for Network Interconnect
              Devices", Informational, RFC 5180, May 2008

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

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

   [RFC5695] Akhter A. "MPLS Forwarding Benchmarking Methodology",
              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

Author's Addresses

   Jan Novak (editor)
   Cisco Systems

Edinburgh,
United Kingdom
Email: janovak@cisco.com

Appendix A: Recommended Report Format
Parameter   Units
----------------------------------- -------------------------------------
Test Case   test case name (section 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 Flow Records per second
   RFC2544 Throughput   packets per second
   (Other RFC2544 Metrics)   (as appropriate)

General Parameters
   Traffic Direction   unidirectional, bidirectional
   DUT Interface Type   Ethernet, POS, ATM, other
   DUT Interface Bandwidth   MegaBits per second

Traffic Specifications
   Number of Traffic Components   (see section 6.4 and 6.5)
   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

Flow monitoring Specifications
   Direction   ingress, egress, both
   Observation Points   DUT interface names
   Cache Size   number of entries
   Active Timeout   seconds
   Inactive 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

Packet Sampling Specifications
   Sampling Method [RFC5475]   systematic, random or none
   Sampling Interval   milliseconds or not applicable
   Sampling Rate   number of packets or not applicable

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

Appendix B: Miscellaneous Tests

   This section lists the tests which could be useful to asses 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 the
      section 6.5, where one traffic component exercises the Flow
      Monitoring Plane and the second traffic component loads only
      Forwarding Plane without affecting Flow monitoring (e.g. it
      creates just one and static Flow Record in the Cache).

      The variance in Flow Monitoring Throughput as 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 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, 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 benchmark test as specified in section 5
      b. One of the results will be how much bandwidth Flow Export
         used on the dedicated Flow Export interface
      c. Change Flow Export configuration to use the test interface
      d. Repeat the benchmark test while the receiver filters out the
         Flow Export data from analysis

      The expected result is that the RFC2544 Throughput achieved in
      step a. is same as the Throughput achieved in step d. 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 mix of packet
   sizes). The packet forwarding tests specified mainly in [RFC2544]
   do not recommend and perform such tests. Flow monitoring is not
   dependent on packet sizes so such a test could be performed during
   the Flow Monitoring Throughput measurement and verify 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 physical bandwidth limitations and changes of base
   forwarding performance with packet size.

B.4 Bursty Traffic

   RFC2544 section 21 discusses and defines the use of bursty
   traffic. It can be used for Flow monitoring testing as well to
   gauge some short term overload DUT capabilities in terms of Flow
   monitoring. The tests benchmark here would not be the Flow
   Expiration Rate the DUT can sustain but the absolute number of
   Flow Records the DUT can process without dropping any single Flow
   Record. The traffic set-up to be used for this test is as follows:

   a. each sent packet creates a new Flow Record
   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
   which is not possible to verify in a black box test manner.

B.5.1 RFC2544 Throughput without Metering Process

    If 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 RFC2544 Throughput with Metering Process

    If only Metering Process is enabled it means that 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 an useful test of the DUT capabilities (this

corresponds to the case when the network operator uses Flow
Monitoring for example for manual denial of service attacks
detection 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 RFC2544 Throughput with Metering and Exporting Process

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

B.6 Tests With Bidirectional Traffic

The test topology on Figure 2 can be expanded to verify Flow
monitoring functionality with bidirectional traffic in two possible
ways:

a. use two sets of interfaces, one for Flow monitoring for ingress
   traffic and one for Flow monitoring egress traffic
b. use exactly same set-up as in Figure 2 but use the interfaces in
   full duplex mode e.g. sending and receiving simultaneously on each
   of them

The set-up in point a. above is in fact equivalent to the set-up with
several Observation Points as already discussed in the section 4.1
and 4.3.1.

For the set-up in point b. same rules should be applied (as per
section 4.1 and 4.3.1) - traffic passing through each Observation
Point SHOULD always create a new Flow Record 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

An additional useful information when analysing the Flow Export data
for the Flow Expiration Rate 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 counting the Flow Records
   and provide the instantaneous (or simply an average over shorter
   time interval than specified in the section 5.4) Flow Export Rate
b. The Flow Export protocol (like IPFIX [RFC5101]) can provide time
   stamps in the Flow Export packets which would allow time based
   analysis and calculate the Flow Export Rate as an average over
   much shorter time interval than specified in the 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.