Terminology for Benchmarking Network-layer Traffic Control Mechanisms

Status of This Memo

Copyright Notice

Abstract

   This document describes terminology for the benchmarking of devices
   that implement traffic control using packet classification based on
   defined criteria. The terminology is to be applied to measurements
   made on the data plane to evaluate IP traffic control mechanisms.
   Rules for packet classification can be based on any field in the IP
   header, such as the Differentiated Services Code Point (DSCP), or any
   field in the packet payload, such as port number.

Table of Contents

1. Introduction

   New terminology is needed because most existing measurements assume
   the absence of congestion and only a single per-hop behavior. This
   document introduces several new terms that will allow measurements to
   be taken during periods of congestion.

   Another key difference from existing terminology is the definition of
   measurements as observed on egress and ingress of a device/system
   under test. Again, the existence of congestion requires the addition
   of egress measurements, as well as of those taken on ingress; without
   observing traffic leaving a device/system, it is not possible to say
   whether traffic-control mechanisms effectively dealt with congestion.

The principal measurements introduced in this document are vectors
for rate, delay, and jitter, all of which can be observed with or
without congestion of the Device Under Test (DUT)/System Under Test
(SUT). This document describes only those terms relevant to
measuring behavior of a DUT or SUT at the egress during periods of
congestion. End-to-end and service-level measurements are beyond the
scope of this document.

2. Existing Definitions

RFC 1224, "Techniques for Managing Asynchronously Generated Alerts"
[St91], is used for 'Time with fine enough units to distinguish
between two events'.

RFC 1242, "Benchmarking Terminology for Network Interconnect
Devices", and RFC 2285, "Benchmarking Terminology for LAN Switching
Devices", should be consulted before attempting to make use of this
document.

RFC 2474, "Definition of the Differentiated Services Field (DS Field)
in the IPv4 and IPv6 Headers", section 2, contains discussions of a
number of terms relevant to network-layer traffic control mechanisms
and should also be consulted.

For the sake of clarity and continuity, this RFC adopts the template
for definitions set out in Section 2 of RFC 1242. Definitions are
indexed and grouped together in sections for ease of reference.

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 BCP 14, RFC 2119
[Br97]. RFC 2119 defines the use of these key words to help make the
intent of standards track documents as clear as possible. While this
document uses these keywords, this document is not a standards track
document.

2.1. Frequently Used Acronyms

```
DA   Destination Address
DS   DiffServ
DSCP DiffServ Code Point
DUT  Device Under Test
IP   Internet Protocol
PHB  Per Hop Behavior
SA   Source Address
SUT  System Under Test
```

3. **Term Definitions**

3.1. **Configuration Terms**

3.1.1. **Classification**

   Definition:
      Selection of packets according to defined rules.

   Discussion:
      Classification determines the per-hop behaviors and traffic
      conditioning functions, such as shaping and dropping, that are to
      be applied to the packet.

      Classification of packets can be based on the DS field or IP
      Precedence in the packet header. Classification can be based on
      other IP header fields, such as IP Source Address (SA),
      Destination Address (DA), and protocol, or on fields in the packet
      payload, such as port number. Classification can also be based on
      ingress interface. It is possible to base classification on
      Multi-Field (MF) criteria such as IP source and destination
      addresses, protocol, and port number. For further discussion of
      packet classification and its network applications, see [Bl98].

   Measurement units:
      n/a

   See Also:
      None

3.1.2. **Codepoint Set**

   Definition:
      The set of all DS Code-points or IP precedence values used during
      the test duration.

   Discussion:
      Describes all the code-point markings associated with packets that
      are input to the DUT/SUT. For each entry in the codepoint set,
      there are associated vectors describing the rate of traffic,
      delay, loss, or jitter containing that particular DSCP or IP
      precedence value.

      The treatment that a packet belonging to a particular code-point
      gets is subject to the DUT classifying packets to map to the
      correct PHB. Moreover, the forwarding treatment in general is
      also dependent on the complete set of offered vectors.

   Measurement Units:
      n/a

   See Also:
      None

3.1.3. Forwarding Congestion

   Definition:
      A condition in which one or more egress interfaces are offered
      more packets than are forwarded.

   Discussion:
      This condition is a superset of the overload definition [Ma98].
      Overload [Ma98] deals with overloading input and output interfaces
      beyond the maximum transmission allowed by the medium. Forwarding
      congestion does not assume ingress interface overload as the only
      source of overload on output interfaces.

      Another difference between Forwarding Congestion and overload
      occurs when the SUT comprises multiple elements, in that
      Forwarding Congestion may occur at multiple points. Consider an
      SUT comprising multiple edge devices exchanging traffic with a
      single core device. Depending on traffic patterns, the edge
      devices may induce Forwarding Congestion on multiple egress
      interfaces on the core device.

      Throughput [Br91] defines the lower boundary of Forwarding
      Congestion. Throughput is the maximum offered rate with no
      Forwarding Congestion. At offered rates above throughput, the
      DUT/SUT is considered to be in a state of Forwarding Congestion.

      Packet Loss, not increased Forwarding Delay, is the external
      observable metric used to indicate the condition of Forwarding
      Congestion. Packet Loss is a deterministic indicator of
      Forwarding Congestion. The condition of increased Forwarding
      Delay without Packet Loss is an indicator of Forwarding Congestion
      known as Incipient Congestion. Incipient Congestion is a non-
      deterministic indicator of Forwarding Congestion [Fl93]. As
      stated in [Ec98], RED [Br98] detects incipient congestion before
      the buffer overflows, but the current Internet environment is
      limited to packet loss as the mechanism for indicating congestion
      to the end-nodes. [Ra99] implies that it is impractical to build
      a black-box test to observe Incipient Congestion. [Ra99] instead
      introduces Explicit Congestion Notification (ECN) as a
      deterministic Black-Box method for observing Incipient Congestion.
      [Ra99] is an Experimental RFC with limited deployment, so ECN is
      not used for this particular methodology. For the purpose of

      "black-box" testing a DUT/SUT, this methodology uses Packet Loss
      as the indicator of Forwarding Congestion.

      Ingress observations alone are not sufficient to cover all cases
      in which Forwarding Congestion may occur. A device with an
      infinite amount of memory could buffer an infinite number of
      packets and eventually forward all of them. However, these
      packets may or may not be forwarded during the test duration.
      Congestion Collapse [Na84] is defined as the state in which
      buffers are full and all arriving packets MUST be dropped across
      the network. Even though ingress interfaces accept all packets
      without loss, Forwarding Congestion is present in this
      hypothetical device.

      The definition presented here explicitly defines Forwarding
      Congestion as an event observable on egress interfaces.
      Regardless of internal architecture, any device exhibiting Packet
      Loss on one or more egress interfaces is experiencing Forwarding
      Congestion.

   Measurement units:
      None

   See Also:
      Gateway Congestion Control Survey [Ma91]

3.1.4. Congestion Management

   Definition:
      An implementation of one or more per-hop behaviors to avoid or
      minimize the condition of congestion.

   Discussion:
      Congestion management may seek either to control congestion or
      avoid it altogether through Classification.

      Congestion avoidance mechanisms seek to prevent congestion before
      it actually occurs.

      Congestion control mechanisms give one or more flows (with a
      discrete IP Precedence or DSCP value) preferential treatment over
      other classes during periods of congestion.

   Measurement units:
      n/a

   See Also:
      Classification

3.1.5. Flow

   Definition:
      A flow is one or more packets sharing a common intended pair of
      ingress and egress interfaces.

   Discussion:
      Packets are grouped by the ingress and egress interfaces they use
      on a given DUT/SUT.

      A flow can contain multiple source IP addresses and/or destination
      IP addresses. All packets in a flow MUST enter on the same
      ingress interface and exit on the same egress interface and have
      some common network layer content.

      Microflows [Ni98] are a subset of flows. As defined in [Ni98],
      microflows require application-to-application measurement. In
      contrast, flows use lower-layer classification criteria. Since
      this document focuses on network-layer classification criteria, it
      concentrates here on the use of network-layer identifiers in
      describing a flow. Flow identifiers also may reside at the data-
      link, transport, or application layers of the OSI model. However,
      identifiers other than those at the network layer are out of scope
      for this document.

      A flow may contain a single code point/IP precedence value or may
      contain multiple values destined for a single egress interface.
      This is determined by the test methodology.

   Measurement units:
      n/a

   See Also:
      Microflow [Ni98]
      Streams

3.2. Measurement Terms

3.2.1. Forwarding Capacity

   Definition:
      The number of packets per second that a device can be observed to
      transmit successfully to the correct egress interface in response
      to a specified offered load while the device drops none of the
      offered packets.

Discussion:
    Forwarding Capacity measures the packet rate at the egress
    interface(s) of the DUT/SUT. In contrast, throughput (as defined
    in RFC 1242) measures the packet rate at the ingress interface(s)
    of the DUT/SUT.

    Ingress-based measurements do not account for queuing of the
    DUT/SUT. Throughput rates can be higher than the Forwarding
    Capacity because of queueing. The difference is dependent upon
    test duration, packet rate, and queue size. Forwarding Capacity,
    as an egress measurement, does take queuing into account.

    Understanding Forwarding Capacity is a necessary precursor to any
    measurement involving Traffic Control Mechanisms. The
    accompanying methodology document MUST take into consideration
    Forwarding Capacity when determining the expected forwarding
    vectors. When the sum of the expected forwarding vectors on an
    interface exceeds the Forwarding Capacity, the Forwarding Capacity
    will govern the forwarding rate.

    This measurement differs from forwarding rate at maximum offered
    load (FRMOL) [Ma98] in that the Forwarding Capacity requires zero
    loss.

Measurement units:
    N-octet packets per second

See Also:
    Throughput [Br91]
    Forwarding Rate at Maximum Offered Load [Ma98]

3.2.2. Conforming Packet

Definition:
    Packets that lie within specific rate, delay, or jitter bounds.

Discussion:
    A DUT/SUT may be configured to allow a given traffic class to
    consume a given amount of bandwidth, or to fall within predefined
    delay or jitter boundaries. All packets that lie within specified
    bounds are then said to be conforming, whereas those outside the
    bounds are nonconforming.

Measurement units:
    n/a

   See Also:
      Expected Vector
      Forwarding Vector
      Offered Vector
      Nonconforming

3.2.3. **Nonconforming Packet**

   Definition:
      Packets that do not lie within specific rate, delay, or jitter
      bounds.

   Discussion:
      A DUT/SUT may be configured to allow a given traffic class to
      consume a given amount of bandwidth, or to fall within predefined
      delay or jitter boundaries. All packets that do not lie within
      these bounds are then said to be nonconforming.

   Measurement units:
      n/a

   See Also:
      Expected Vector
      Forwarding Vector
      Offered Vector
      Conforming

3.2.4. **Forwarding Delay**

   Definition:
      The time interval starting when the last bit of the input IP
      packet is offered to the input port of the DUT/SUT and ending when
      the last bit of the output IP packet is received from the output
      port of the DUT/SUT.

   Discussion:
      The delay time interval MUST be externally observed. The delay
      measurement MUST NOT include delays added by test bed components
      other than the DUT/SUT, such as propagation time introduced by
      cabling or non-zero delay added by the test instrument.
      Forwarding Delay differs from latency [Br91] and one-way delay
      [Al99] in several key regards:

      1. Latency [Br91] assumes knowledge of whether the DUT/SUT uses
         "store and forward" or "bit forwarding" technology. Forwarding
         Delay is the same metric, measured the same way, regardless of
         the architecture of the DUT/SUT.

   2. Forwarding Delay is a last-in, last-out (LILO) measurement,
      unlike the last-in, first-out method [Br91] or the first-in,
      last-out method [Al99].

      The LILO method most closely simulates the way a network-layer
      device actually processes an IP datagram. IP datagrams are not
      passed up and down the stack unless they are complete, and
      processing begins only once the last bit of the IP datagram has
      been received.

      Further, the LILO method has an additive property, where the
      sum of the parts MUST equal the whole. This is a key
      difference from [Br91] and [Al99]. For example, the delay
      added by two DUTs MUST equal the sum of the delay of the DUTs.
      This may or may not be the case with [Br91] and [Al99].

   3. Forwarding Delay measures the IP datagram only, unlike [Br91],
      which also includes link-layer overhead.

      A metric focused exclusively on the Internet protocol relieves
      the tester from specifying the start/end for every link-layer
      protocol that IP runs on. This avoids the need to determine
      whether the start/stop delimiters are included. It also allows
      the use of heterogeneous link-layer protocols in a test.

   4. Forwarding Delay can be measured at any offered load, whereas
      the latency methodology [Br99] recommends measurement at, and
      only at, the throughput level. Comparing the Forwarding Delay
      below the throughput to Forwarding Delay above the Forwarding
      Capacity will give insight to the traffic control mechanisms.

      For example, non-congested delay may be measured with an
      offered load that does not exceed the Forwarding Capacity,
      while congested delay may involve an offered load that exceeds
      the Forwarding Capacity.

      Note: Forwarding Delay SHOULD NOT be used as an absolute
      indicator of DUT/SUT Forwarding Congestion. While Forwarding
      Delay may rise when offered load nears or exceeds the
      Forwarding Capacity, there is no universal point at which
      Forwarding Delay can be said to indicate the presence or
      absence of Forwarding Congestion.

   Measurement units:
      milliseconds

   See Also:
      Latency [Br91]
      Latency [Al99]
      One-way Delay [Br99]

3.2.5. Jitter

   Definition:
      The absolute value of the difference between the Forwarding Delay
      of two consecutive received packets belonging to the same stream.

   Discussion:
      The Forwarding Delay fluctuation between two consecutive received
      packets in a stream is reported as the jitter. Jitter can be
      expressed as $|D(i) - D(i-1)|$, where D equals the Forwarding Delay
      and i is the order the packets were received.

      Under loss, jitter can be measured between non-consecutive test
      sequence numbers. When IP Traffic Control Mechanisms are dropping
      packets, fluctuating Forwarding Delay may be observed. Jitter
      MUST be able to benchmark the delay variation independently of
      packet loss.

      Jitter is related to the IPDV [De02] (IP Delay Variation) by
      taking the absolute value of the ipdv. The two metrics will
      produce different mean values. Mean Jitter will produce a
      positive value, where the mean ipdv is typically zero. Also, IPDV
      is undefined when one packet from a pair is lost.

   Measurement units:
      milliseconds

   See Also:
      Forwarding Delay
      Jitter variation [Ja99]
      ipdv [De02]
      interarrival jitter [Sc96]

3.2.6. Undifferentiated Response

   Definition:
      The vector(s) obtained when mechanisms used to support diff-serv
      or IP precedence are disabled.

   Discussion:
      Enabling diff-serv or IP precedence mechanisms may impose
      additional processing overhead for packets. This overhead may
      degrade performance even when traffic belonging to only one class,

      the best-effort class, is offered to the device. Measurements
      with "undifferentiated response" SHOULD be made to establish a
      baseline.

      The vector(s) obtained with DSCP or IP precedence enabled can be
      compared to the undifferentiated response to determine the effect
      of differentiating traffic.

   Measurement units:
      n/a

3.3. Sequence Tracking

3.3.1. Test Sequence Number

   Definition:
      A field in the IP payload portion of the packet that is used to
      verify the order of the packets on the egress of the DUT/SUT.

   Discussion:
      The traffic generator sets the test sequence number value. Upon
      receipt of the packet, the traffic receiver checks the value.
      The traffic generator changes the value on each packet transmitted
      based on an algorithm agreed to by the traffic receiver.

      The traffic receiver keeps track of the sequence numbers on a
      per-stream basis. In addition to the number of received packets,
      the traffic receiver may also report the number of in-sequence
      packets, the number of out-of-sequence packets, the number of
      duplicate packets, and the number of reordered packets. The
      RECOMMENDED algorithm to change the sequence number on sequential
      packets is an incrementing value.

   Measurement units:
      n/a

   See Also:
      Stream

3.3.2. Stream

   Definition:
      A group of packets tracked as a single entity by the traffic
      receiver. A stream MUST share common content, such as type (IP,
      UDP), IP SA/DA, packet size, or payload.

   Discussion:
      Streams are tracked by test sequence number or "unique signature
      field" [Ma00]. Streams define how individual packet statistics
      are grouped together to form an intelligible summary.

      Common stream groupings would be by egress interface, destination
      address, source address, DSCP, or IP precedence. A stream using
      test sequence numbers can track the ordering of packets as they
      traverse the DUT/SUT.

      Streams are not restricted to a pair of source and destination
      interfaces as long as all packets are tracked as a single entity.
      A multicast stream can be forwarded to multiple destination
      interfaces.

   Measurement units:
      n/a

   See Also:
      Flow
      Microflow [Ni98]
      Test sequence number

3.3.3. In-Sequence Packet

   Definition:
      A received packet with the expected Test Sequence number.

   Discussion:
      In-sequence is done on a stream level. As packets are received on
      a stream, each packet's Test Sequence number is compared with the
      previous packet. Only packets that match the expected Test
      Sequence number are considered in-sequence.

      Packets that do not match the expected Test Sequence number are
      counted as "not in-sequence" or out-of-sequence. Every packet
      that is received is either in-sequence or out-of-sequence.
      Subtracting the in-sequence from the received packets (for that
      stream), the tester can derive the out-of-sequence count.

      Two types of events will prevent the in-sequence from
      incrementing: packet loss and reordered packets.

   Measurement units:
      Packet count

   See Also:
      Stream
      Test Sequence number

3.3.4. Out-of-Order Packet

   Definition:
      A received packet with a sequence number less than the sequence
      number of any previously arriving packet.

   Discussion:
      As a stream of packets enters a DUT/SUT, they include a Stream
      Test Sequence number indicating the order the packets were sent to
      the DUT/SUT. On exiting the DUT/SUT, these packets may arrive in
      a different order. Each packet that was reordered is counted as
      an Out-of-Order Packet.

      Certain streaming protocols (such as TCP) require the packets to
      be in a certain order. Packets outside this are dropped by the
      streaming protocols even though they were properly received by the
      IP layer. The type of reordering tolerated by a streaming
      protocol varies from protocol to protocol, and also by
      implementation.

      Packet loss does not affect the Out-of-Order Packet count. The
      Out-of-Order Packet count is impacted only by packets that were
      not received in the order that they were transmitted.

   Measurement units:
      packets

   See Also:
      Stream
      Test Sequence number
      Packet Reordering Metric for IPPM [Mo03]

3.3.5. Duplicate Packet

   Definition:
      A received packet with a Test Sequence number matching a
      previously received packet.

   Discussion:
      A Duplicate Packet is a packet that the DUT/SUT has successfully
      transmitted out an egress interface more than once. The egress
      interface has previously forwarded this packet.

   A Duplicate Packet SHOULD be a bit-for-bit copy of an already
   transmitted packet (including Test Sequence number). If the
   Duplicate Packet traversed different paths through the DUT/SUT,
   some fields (such as TTL or checksum) may have changed.

   A multicast packet is not a Duplicate Packet by definition. For a
   given IP multicast group, a DUT/SUT SHOULD forward a packet once
   on a given egress interface provided the path to one or more
   multicast receivers is through that interface. Several egress
   interfaces will transmit the same packet, but only once per
   interface.

   To detect a Duplicate Packet, each packet offered to the DUT/SUT
   MUST contain a unique packet-by-packet identifier.

Measurement units:
   Packet count

See Also:
   Stream
   Test Sequence number

3.4. Vectors

A vector is a group of packets all matching a specific
classification criteria, such as DSCP. Vectors are
identified by the classification criteria and benchmarking
metrics, such as a Forwarding Capacity, Forwarding Delay,
or Jitter.

3.4.1. Intended Vector

Definition:
   A description of the configuration on an external source
   for the attempted rate of a stream transmitted to a DUT/SUT
   matching specific classification rules.

Discussion:
   The Intended Vector of a stream influences the benchmark
   measurements. The Intended Vector is described by the
   classification criteria and attempted rate.

Measurement Units:
   N-bytes packets per second

See Also:
   Stream
   Offered Vector
   Forwarding Vector

3.4.2. Offered Vector

Definition:
   A description for the attempted rate of a stream offered to
   a DUT/SUT matching specific classification rules.

Discussion:
   The Offered Vector of a stream influences the benchmark
   measurements. The Offered Vector is described by the
   classification criteria and offered rate.

Measurement Units:
   N-bytes packets per second

See Also:
   Stream
   Intended Vector
   Forwarding Vector

3.4.3. Expected Vectors

3.4.3.1. Expected Forwarding Vector

Definition:
   A description of the expected output rate of packets matching a
   specific classification, such as DSCP.

Discussion:
   The value of the Expected Forwarding Vector is dependent on the
   set of offered vectors and Classification configuration on the
   DUT/SUT. The DUT is configured in a certain way so that
   classification occurs when a traffic mix consisting of multiple
   streams is applied.

   This term captures the expected forwarding behavior from the DUT
   receiving multiple Offered Vectors. The actual algorithm or
   mechanism the DUT uses to achieve service differentiation is
   implementation specific and is not important when describing the
   Expected Forwarding Vector.

Measurement units:
   N-octet packets per second

   See Also:
      Classification
      Stream
      Intended Vector
      Offered Vector

3.4.3.2. Expected Loss Vector

   Definition:
      A description of the percentage of packets having a specific
      classification that should not be forwarded.

   Discussion:
      The value of the Expected Loss Vector is dependent on the set of
      offered vectors and Classification configuration on the DUT/SUT.
      The DUT is configured in a certain way so that classification
      occurs when a traffic mix consisting of multiple streams is
      applied.

      This term captures the expected forwarding behavior from the DUT
      receiving multiple Offered Vectors. The actual algorithm or
      mechanism the DUT uses to achieve service differentiation is
      implementation specific and is not important when describing the
      Expected Loss Vector.

   Measurement Units:
      Percentage of intended packets expected to be dropped.

   See Also:
      Classification
      Stream
      Intended Vector
      Offered Vector
      One-way Packet Loss Metric [Ka99]

3.4.3.3. Expected Sequence Vector

   Definition:
      A description of the expected in-sequence packets matching a
      specific classification, such as DSCP.

   Discussion:
      The value of the Expected Sequence Vector is dependent on the set
      of offered vectors and Classification configuration on the
      DUT/SUT. The DUT is configured in a certain way so that
      classification occurs when a traffic mix consisting of multiple
      streams is applied.

      This term captures the expected forwarding behavior from the DUT
      receiving multiple Offered Vectors. The actual algorithm or
      mechanism the DUT uses to achieve service differentiation is
      implementation specific and is not important when describing the
      Expected Sequence Vector.

   Measurement Units:
      N-octet packets per second

   See Also:
      Classification
      Stream
      In-Sequence Packet
      Intended Vector
      Offered Vector

3.4.3.4. Expected Delay Vector

   Definition:
      A description of the expected instantaneous Forwarding Delay for
      packets matching a specific classification, such as DSCP.

   Discussion:
      The value of the Expected Delay Vector is dependent on the set of
      offered vectors and Classification configuration on the DUT/SUT.
      The DUT is configured in a certain way so that classification
      occurs when a traffic mix consisting of multiple streams is
      applied.

      This term captures the expected forwarding behavior from the DUT
      receiving multiple Offered Vectors. The actual algorithm or
      mechanism the DUT uses to achieve service differentiation is
      implementation specific and is not important when describing the
      Expected Delay Vector.

   Measurement units:
      milliseconds

   See Also:
      Classification
      Stream
      Forwarding Delay
      Intended Vector
      Offered Vector

3.4.3.5. **Expected Average Delay Vector**

   Definition:
      A description of the expected average Forwarding Delay for packets
      matching a specific classification, such as DSCP.

   Discussion:
      The value of the Expected Average Delay Vector is dependent on the
      set of offered vectors and Classification configuration on the
      DUT/SUT. The DUT is configured in a certain way so that
      classification occurs when a traffic mix consisting of multiple
      streams is applied.

      This term captures the expected forwarding behavior from the DUT
      receiving multiple Offered Vectors. The actual algorithm or
      mechanism the DUT uses to achieve service differentiation is
      implementation specific and is not important when describing the
      Expected Average Delay Vector.

   Measurement units:
      milliseconds

   See Also:
      Classification
      Stream
      Forwarding Delay
      Intended Vector
      Offered Vector
      Expected Delay Vector

3.4.3.6. **Expected Maximum Delay Vector**

   Definition:
      A description of the expected maximum Forwarding Delay for packets
      matching a specific classification, such as DSCP.

   Discussion:
      The value of the Expected Maximum Delay Vector is dependent on the
      set of offered vectors and Classification configuration on the
      DUT/SUT. The DUT is configured in a certain way so that
      classification occurs when a traffic mix consisting of multiple
      streams is applied.

      This term captures the expected forwarding behavior from the DUT
      receiving multiple Offered Vectors. The actual algorithm or
      mechanism the DUT uses to achieve service differentiation is
      implementation specific and is not important when describing the
      Expected Maximum Delay Vector.

   Measurement units:
      milliseconds

   See Also:
      Classification
      Stream
      Forwarding Delay
      Intended Vector
      Offered Vector
      Expected Delay Vector

3.4.3.7. **Expected Minimum Delay Vector**

   Definition:
      A description of the expected minimum Forwarding Delay for packets
      matching a specific classification, such as DSCP.

   Discussion:
      The value of the Expected Minimum Delay Vector is dependent on the
      set of offered vectors and Classification configuration on the
      DUT/SUT. The DUT is configured in a certain way so that
      classification occurs when a traffic mix consisting of multiple
      streams is applied.

      This term captures the expected forwarding behavior from the DUT
      receiving multiple Offered Vectors. The actual algorithm or
      mechanism the DUT uses to achieve service differentiation is
      implementation specific and is not important when describing the
      Expected Minimum Delay Vector.

   Measurement units:
      milliseconds

   See Also:
      Classification
      Stream
      Forwarding Delay
      Intended Vector
      Offered Vector
      Expected Delay Vector

3.4.3.8. **Expected Instantaneous Jitter Vector**

   Definition:
      A description of the expected Instantaneous Jitter between two
      consecutive packets arrival times matching a specific
      classification, such as DSCP.

Discussion:
   Instantaneous Jitter is the absolute value of the difference
   between the Forwarding Delay measurement of two packets belonging
   to the same stream.

   The Forwarding Delay fluctuation between two consecutive packets
   in a stream is reported as the "Instantaneous Jitter".
   Instantaneous Jitter can be expressed as $|D(i) - D(i-1)|$, where D
   equals the Forwarding Delay and i is the test sequence number.
   Packets lost are not counted in the measurement.

   The Forwarding Vector may contain several Jitter Vectors. For n
   packets received in a Forwarding Vector, there is a total of (n-1)
   Instantaneous Jitter Vectors.

Measurement units:
   milliseconds

See Also:
   Classification
   Stream
   Jitter
   Intended Vector
   Offered Vector

3.4.3.9. **Expected Average Jitter Vector**

Definition:
   A description of the expected average jitter for packets arriving
   in a stream matching a specific classification, such as DSCP.

Discussion:
   Average Jitter Vector is the average of all the Instantaneous
   Jitter Vectors measured during the test duration for the same
   stream.

   The value of the Expected Average Jitter Vector is dependent on
   the set of offered vectors and Classification configuration on the
   DUT/SUT. The DUT is configured in a certain way so that
   classification occurs when a traffic mix consisting of multiple
   streams is applied.

   This term captures the expected forwarding behavior from the DUT
   receiving multiple Offered Vectors. The actual algorithm or
   mechanism the DUT uses to achieve service differentiation is
   implementation specific and is not important when describing the
   Expected Average Jitter Vector.

   Measurement units:
      milliseconds

   See Also:
      Classification
      Stream
      Jitter
      Intended Vector
      Offered Vector
      Expected Instantaneous Jitter Vector

3.4.3.10. **Expected Peak-to-peak Jitter Vector**

   Definition:
      A description of the expected maximum variation in the Forwarding
      Delay of packet arrival times for packets arriving in a stream
      matching a specific classification, such as DSCP.

   Discussion:
      Peak-to-peak Jitter Vector is the maximum Forwarding Delay minus
      the minimum Forwarding Delay of the packets (in a vector)
      forwarded by the DUT/SUT.

      Peak-to-peak Jitter is not derived from the Instantaneous Jitter
      Vector. Peak-to-peak Jitter is based upon all the packets during
      the test duration, not just two consecutive packets.

      The value of the Expected Peak-to-peak Jitter Vector is dependent
      on the set of offered vectors and Classification configuration on
      the DUT/SUT. The DUT is configured in a certain way so that
      classification occurs when a traffic mix consisting of multiple
      streams is applied.

      This term captures the expected forwarding behavior from the DUT
      receiving multiple Offered Vectors. The actual algorithm or
      mechanism the DUT uses to achieve service differentiation is
      implementation specific and is not important when describing the
      Expected Peak-to-peak Jitter Vector.

   Measurement units:
      milliseconds

   See Also:
      Classification
      Stream
      Jitter
      Intended Vector
      Offered Vector
      Expected Instantaneous Jitter Vector
      Expected Average Jitter Vector

3.4.4. **Output Vectors**

3.4.4.1. **Forwarding Vector**

   Definition:
      The number of packets per second for a stream matching a specific
      classification, such as DSCP, that a DUT/SUT is measured to
      forward to the correct destination interface successfully in
      response to an offered vector.

   Discussion:
      Forwarding Vector is expressed as a combination of values: the
      classification rules AND the measured packets per second for the
      stream matching the classification rules. Forwarding Vector is a
      per-hop measurement. The DUT/SUT MAY remark the specific DSCP (or
      IP precedence) value for a multi-hop measurement. The stream
      remains the same.

   Measurement units:
      N-octet packets per second

   See Also:
      Classification
      Stream
      Forwarding Capacity
      Intended Vector
      Offered Vector
      Expected Vector

3.4.4.2. **Loss Vector**

   Definition:
      The percentage of packets per second for a stream matching a
      specific classification, such as DSCP, that a DUT/SUT is measured
      not to transmit to the correct destination interface in response
      to an offered vector.

   Discussion:
      Loss Vector is expressed as a combination of values: the
      classification rules AND the measured percentage value of packet
      loss. Loss Vector is a per-hop measurement. The DUT/SUT MAY
      remark the specific DSCP or IP precedence value for a multi-hop
      measurement. The stream remains the same.

   Measurement Units:
      Percentage of packets

   See Also:
      Classification
      Stream
      Intended Vector
      Offered Vector
      Expected Vector
      One-way Packet Loss Metric [Ka99]

3.4.4.3. Sequence Vector

   Definition:
      The number of packets per second for all packets in a stream
      matching a specific classification, such as DSCP, that a DUT/SUT
      is measured to transmit in sequence to the correct destination
      interface in response to an offered vector.

   Discussion:
      Sequence Vector is expressed as a combination of values: the
      classification rules AND the number of packets per second that are
      in-sequence.

      Sequence Vector is a per-hop measurement. The DUT/SUT MAY remark
      the specific DSCP or IP precedence value for a multi-hop
      measurement. The stream remains the same.

   Measurement Units:
      N-octet packets per second

   See Also:
      Classification
      Stream
      In-sequence Packet
      Intended Vector
      Offered Vector
      Expected Vector

3.4.4.4. **Instantaneous Delay Vector**

   Definition:
      The instantaneous Forwarding Delay for a packet in a stream
      matching a specific classification, such as DSCP, that a DUT/SUT
      is measured to transmit to the correct destination interface
      successfully in response to an offered vector.

   Discussion:
      Instantaneous Delay Vector is expressed as a combination of
      values: the classification rules AND Forwarding Delay. For every
      packet received in a Forwarding Vector, there is a corresponding
      Instantaneous Delay Vector.

      Instantaneous Delay Vector is a per-hop measurement. The DUT/SUT
      MAY remark the specific DSCP or IP precedence value for a multi-
      hop measurement. The stream remains the same.

      Instantaneous Delay Vector can be obtained at any offered load.
      It is RECOMMENDED that this vector be obtained at or below the
      Forwarding Capacity in the absence of Forwarding Congestion. For
      congested Forwarding Delay, run the offered load above the
      Forwarding Capacity.

   Measurement Units:
      milliseconds

   See Also:
      Classification
      Stream
      Forwarding Capacity
      Forwarding Delay
      Intended Vector
      Offered Vector
      Expected Delay Vector

3.4.4.5. **Average Delay Vector**

   Definition:
      The average Forwarding Delay for packets in a stream matching a
      specific classification, such as DSCP, that a DUT/SUT is measured
      to transmit to the correct destination interface successfully in
      response to an offered vector.

   Discussion:
      Average Delay Vector is expressed as combination of values: the
      classification rules AND average Forwarding Delay.

      The average Forwarding Delay is computed by averaging all the
      Instantaneous Delay Vectors for a given stream.

      Average Delay Vector is a per-hop measurement. The DUT/SUT MAY
      remark the specific DSCP or IP precedence value for a multi-hop
      measurement. The stream remains the same.

      Average Delay Vector can be obtained at any offered load. It is
      recommended that the offered load be at or below the Forwarding
      Capacity in the absence of congestion. For congested Forwarding
      Delay, run the offered load above the Forwarding Capacity.

   Measurement Units:
      milliseconds

   See Also:
      Classification
      Stream
      Forwarding Capacity
      Forwarding Delay
      Intended Vector
      Offered Vector
      Expected Delay Vector
      Instantaneous Delay Vector

3.4.4.6. Maximum Delay Vector

   Definition:
      The maximum Forwarding Delay for packets in a stream matching a
      specific classification, such as DSCP, that a DUT/SUT is measured
      to transmit to the correct destination interface successfully in
      response to an offered vector.

   Discussion:
      Maximum Delay Vector is expressed as combination of values: the
      classification rules AND maximum Forwarding Delay.

      The maximum Forwarding Delay is computed by selecting the highest
      value from the Instantaneous Delay Vectors for a given stream.

      Maximum Delay Vector is a per-hop measurement. The DUT/SUT MAY
      remark the specific DSCP or IP precedence value for a multi-hop
      measurement. The stream remains the same.

      Maximum Delay Vector can be obtained at any offered load. It is
      recommended that the offered load be at or below the Forwarding
      Capacity in the absence of congestion. For congested Forwarding
      Delay, run the offered load above the Forwarding Capacity.

   Measurement Units:
      milliseconds

   See Also:
      Classification
      Stream
      Forwarding Capacity
      Forwarding Delay
      Intended Vector
      Offered Vector
      Expected Delay Vector
      Instantaneous Delay Vector

3.4.4.7. **Minimum Delay Vector**

   Definition:
      The minimum Forwarding Delay for packets in a stream matching a
      specific classification, such as DSCP, that a DUT/SUT is measured
      to transmit to the correct destination interface successfully in
      response to an offered vector.

   Discussion:
      Minimum Delay Vector is expressed as a combination of values: the
      classification rules AND minimum Forwarding Delay. The minimum
      Forwarding Delay is computed by selecting the lowest value from
      the Instantaneous Delay Vectors for a given stream.

      Minimum Delay Vector is a per-hop measurement. The DUT/SUT MAY
      remark the specific DSCP or IP precedence value for a multi-hop
      measurement. The stream remains the same.

      Minimum Delay Vector can be obtained at any offered load. It is
      recommended that the offered load be at or below the Forwarding
      Capacity in the absence of congestion. For congested Forwarding
      Delay, run the offered load above the Forwarding Capacity.

   Measurement Units:
      milliseconds

   See Also:
      Classification
      Stream
      Forwarding Capacity
      Forwarding Delay
      Intended Vector
      Offered Vector
      Expected Delay Vector

3.4.4.8. **Instantaneous Jitter Vector**

   Definition:
      The jitter for two consecutive packets in a stream matching a
      specific classification, such as DSCP, that a DUT/SUT is measured
      to transmit to the correct destination interface successfully in
      response to an offered vector.

   Discussion:
      Instantaneous Jitter is the absolute value of the difference
      between the Forwarding Delay measurement of two packets belonging
      to the same stream.

      The Instantaneous Jitter vector is expressed as a pair of numbers.
      Both the specific DSCP (or IP precedence) value AND jitter value
      combine to make a vector.

      The Forwarding Delay fluctuation between two consecutive packets
      in a stream is reported as the "Instantaneous Jitter".
      Instantaneous Jitter Vector can be expressed as $|D(i) - D(i-1)|$,
      where D equals the Forwarding Delay and i is the test sequence
      number. Packets lost are not counted in the measurement.

      The Instantaneous Jitter Vector is a per-hop measurement. The
      DUT/SUT MAY remark the specific DSCP or IP precedence value for a
      multi-hop measurement. The stream remains the same.

      There may be several Instantaneous Jitter Vectors for a single
      stream. For n packets measured, there may be (n-1) Instantaneous
      Jitter Vectors.

   Measurement units:
      milliseconds

See Also:
   Classification
   Stream
   Forwarding Delay
   Jitter
   Forwarding Vector
   Expected Vectors

3.4.4.9. **Average Jitter Vector**

   Definition:
      The average jitter for packets in a stream matching a specific
      classification, such as DSCP, that a DUT/SUT is measured to
      transmit to the correct destination interface successfully in
      response to an offered vector.

   Discussion:
      Average jitter is calculated by the average of all the
      Instantaneous Jitter Vectors of the same stream measured during
      the test duration. Average Jitter Vector is expressed as a
      combination of values: the classification rules AND average
      Jitter.

      Average Jitter Vector is a per-hop measurement. The DUT/SUT MAY
      remark the specific DSCP or IP precedence value for a multi-hop
      measurement. The stream remains the same.

   Measurement units:
      milliseconds

   See Also:
      Classification
      Stream
      Jitter
      Forwarding Vector
      Expected Vector
      Instantaneous Jitter Vector

3.4.4.10. **Peak-to-peak Jitter Vector**

   Definition:
      The maximum possible variation in the Forwarding Delay for packets
      in a stream matching a specific classification, such as DSCP, that
      a DUT/SUT is measured to transmit to the correct destination
      interface successfully in response to an offered vector.

Discussion:
   Peak-to-peak Jitter Vector is calculated by subtracting the
   maximum Forwarding Delay from the minimum Forwarding Delay of the
   packets forwarded by the DUT/SUT. Jitter vector is expressed as a
   combination of values: the classification rules AND peak-to-peak
   Jitter.

   Peak-to-peak Jitter is not derived from the Instantaneous Jitter
   Vector. Peak-to-peak Jitter is based upon all the packets during
   the test duration, not just two consecutive packets.

Measurement units:
   milliseconds

See Also:
   Jitter
   Forwarding Vector
   Stream
   Expected Vectors
   Instantaneous Jitter Vector
   Average Jitter Vector

4. Security Considerations

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

Packets with unintended and/or unauthorized DSCP or IP precedence
values may present security issues. Determining the security
consequences of such packets is out of scope for this document.

5. Acknowledgements

The authors gratefully acknowledge the contributions of the IETF's
Benchmarking Methodology Working Group members in reviewing this
document. The authors would like to express our thanks to David
Newman for his consistent and valuable assistance throughout the
development of this document. The authors would also like to thank
Al Morton and Kevin Dubray for their ideas and support.

6. References

6.1. Normative References

[Br91] Bradner, S., "Benchmarking terminology for network
       interconnection devices", RFC 1242, July 1991.

[Br97] Bradner, S., "Key words for use in RFCs to Indicate
       Requirement Levels", BCP 14, RFC 2119, March 1997.

[Br98] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S.,
       Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge,
       C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski,
       J., and L. Zhang, "Recommendations on Queue Management and
       Congestion Avoidance in the Internet", RFC 2309, April 1998.

[Ma98] Mandeville, R., "Benchmarking Terminology for LAN Switching
       Devices", RFC 2285, February 1998.

[Ni98] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition
       of the Differentiated Services Field (DS Field) in the IPv4
       and IPv6 Headers", RFC 2474, December 1998.

[St91] Steinberg, L., "Techniques for managing asynchronously
       generated alerts", RFC 1224, May 1991.

6.2. Informative References

[Al99] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay
       Metric for IPPM", RFC 2679, September 1999.

[Bl98] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and
       W. Weiss, "An Architecture for Differentiated Service", RFC
       2475, December 1998.

[Br99] Bradner, S. and J. McQuaid, "Benchmarking Methodology for
       Network Interconnect Devices", RFC 2544, March 1999.

[De02] Demichelis, C. and P. Chimento, "IP Packet Delay Variation
       Metric for IP Performance Metrics (IPPM)", RFC 3393, November
       2002.

[Ec98] http://www3.ietf.org/proceedings/98mar/98mar-edited-135.htm

[Fl93] Floyd, S., and Jacobson, V., "Random Early Detection gateways
       for Congestion Avoidance", IEEE/ACM Transactions on
       Networking, V.1 N.4, August 1993, p. 397-413. URL
       "ftp://ftp.ee.lbl.gov/papers/early.pdf".

   [Ja99] Davie, B., Charny, A., Bennet, J.C., Benson, K., Le Boudec,
          J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis,
          "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246,
          March 2002.

   [Ka99] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet
          Loss Metric for IPPM", RFC 2680, September 1999.

   [Ma91] Mankin, A. and K. Ramakrishnan, "Gateway Congestion Control
          Survey", RFC 1254, August 1991.

   [Ma00] Mandeville, R. and J. Perser, "Benchmarking Methodology for
          LAN Switching Devices", RFC 2889, August 2000.

   [Mo03] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S.,
          Perser, J., "Packet Reordering Metric for IPPM", Work in
          Progress.

   [Na84] Nagle, J., "Congestion control in IP/TCP internetworks", RFC
          896, January 1984.

   [Ra99] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of
          Explicit Congestion Notification (ECN) to IP", RFC 3168,
          September 2001.

   [Sc96] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson,
          "RTP: A Transport Protocol for Real-Time Applications", STD
          64, RFC 3550, July 2003.

Authors' Addresses

   Jerry Perser
   Veriwave
   8770 SW Nimbus Ave.
   Suite B
   Beaverton, OR 97008   USA
   USA

   Phone: + 1 818 338 4112
   EMail: jerry@perser.org

   Scott Poretsky
   Reef Point Systems
   8 New England Executive Park
   Burlington, MA 01803
   USA

   Phone: + 1 508 439 9008
   EMail: sporetsky@reefpoint.com

   Shobha Erramilli
   Telcordia Technologies
   331 Newman Springs Road
   Red Bank, New Jersey 07701
   USA

   EMail: shobha@research.telcordia.com

   Sumit Khurana
   Motorola
   7700 West Parmer Ln.
   Austin, TX 78729
   USA

   Phone: +1 512 996 6604
   Email: skhurana@motorola.com