Network Working Group
INTERNET-DRAFT
Expires in: April 2003

Jerry Perser
Spirent
David Newman
Network Test
Sumit Khurana
Telcordia
Shobha Erramilli
QNetworx
Scott Poretsky
Avici Systems
October 2002

Terminology for Benchmarking Network-layer Traffic Control Mechanisms

<draft-ietf-bmwg-dsmterm-04.txt>

Status of this Memo

Table of Contents

3.2 Measurement Terms

1. Introduction

This document describes terminology for the benchmarking of devices that implement traffic control based on IP precedence or diff-serv code point criteria.

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 as well as ingress of a device/system under test. Again, the existence of congestion requires the addition of egress measurements as well as 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 DUT/SUT.

This document describes only those terms relevant to measuring behavior of a device or a group of devices using one of these two mechanisms. End-to-end and service-level measurements are beyond the scope of this document.

2. Existing definitions

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 RFC 2119.

3. Term definitions

3.1 Configuration Terms

3.1.1 Classification

Definition:

Selection of packets based on the contents of packet header according to defined rules.

Discussion:

Packets can be selected based on the DS field or IP Precedence in the packet header. Classification can also be based on Multi-Field (MF) criteria such as IP Source and destination addresses, protocol and port number.

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

Measurement units:

n/a

See Also:

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:

3.1.3 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 [2]. The overload definition assumes the congestion is introduced strictly by the tester on ingress of a DUT/SUT. That may or may not be the case here.

Another difference between congestion and overload occurs when the SUT comprises multiple elements, in that 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 congestion on multiple egress interfaces on the core device. In contrast, overload [1] deals only with overload on ingress.

Throughput [1] defines the lower boundary of congestion. Throughput is the maximum offered rate with no congestion. At offered rates above throughput, the DUT/SUT is considered congested.

Ingress observations alone are not sufficient to cover all cases in which congestion may occur. A device with an infinite amount of memory could buffer an infinite amount of packets, and eventually forward all of them. However, these packets may or may not be forwarded during the test duration. Even though ingress interfaces accept all packets without loss, this hypothetical device may still be congested.

The definition presented here explicitly defines congestion as an event observable on egress interfaces. Regardless of internal architecture, any device that cannot forward packets on one or more egress interfaces is congested.

Measurement units:
none

See Also:
Gateway Congestion Control Survey [8]

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. Such mechanisms classify packets based upon IP Precedence or DSCP settings in a packet's IP

header.

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

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

Measurement units:
n/a

See Also:

3.1.5 Flow

Definition:
A flow is a one or more of packets sharing a common intended pair of source and destination 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 [3] are a subset of flows. As defined in [3], microflows require application-to-application measurement. In contrast, flows use lower-layer classification criteria. Since this document focuses on network-layer classification criteria, we concentrate 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 ISO 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 [3]
Streams

3.2 Measurement Terms

3.2.1 Channel Capacity

Definition:
The maximum forwarding rate [2] at which none of the offered packets are dropped by the DUT/SUT.

Discussion:
Channel 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 is based on the ingress interface(s) of the DUT/SUT.

Ingress-based measurements do not account for congestion of the DUT/SUT. Channel capacity, as an egress measurement, does take congestion into account.

Understanding channel capacity is a necessary precursor to any measurement involving congestion. Throughput numbers can be higher than channel capacity because of queueing.

This measurement differs from forwarding rate at maximum offered load (FRMOL) [2] in that it is intolerant of loss.

Measurement units:
N-octet packets per second

See Also:
Throughput [1]
Forwarding Rate at Maximum Offered Load [2]

3.2.2 Conforming

Definition:
Packets which 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

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 Delay

Definition:
The time interval starting when the last bit of the input IP packet reaches the input port of the DUT/SUT and ending when the last bit of the output IP packet is seen on the output port of the DUT/SUT.

Discussion:
Delay is measured the same regardless of the type of DUT/SUT. Latency [1] require some knowledge of whether the DUT/SUT is a "store and forward" or a "bit forwarding" device. The fact

that a DUT/SUT's technology has a lower delay than another technology should be visible.

By specifying the metric to be inside the Internet protocol, the tester is relieved from specifying the start/end for every data link layer protocol that IP runs on. This avoids determining if the start/end delimiters are included in the frame. Also heterogeneous data link protocol can be used in a test.

The measurement point at the end closely simulates the way an internet datagram is processed. An internet datagram is not passed up or down the stack unless it is complete. Completion occurs once the last bit of the IP packet has been received.

Delay can be run at any offered load. Recommend at or below the channel capacity for non-congested delay. For congested delay, run the offered load above the channel capacity.

Measurement units:
Seconds.

See Also:
Latency [1]

3.2.5 Jitter

Definition:
The absolute value of the difference between the arrival delay of two consecutive packets belonging to the same stream.

Discussion:
The delay fluctuation between two consecutive packets in a stream is reported as the jitter. Jitter can be expressed as |D(i) - D(i-1)| where D equals the delay and i is the test sequence number. The measurement does not include lost packets.

Jitter can be determined by the ipdv [6] (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.

Measurement units:
Seconds

See Also:

Jitter variation [5]
ipdv [6]
interarrival jitter [7]

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 DSCPs 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 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 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) 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

3.3.2 Out-of-order Packet

Definition:
A received packet with a Test Sequence number less that expected.

Discussion:
As a stream of packets enter 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 re-ordered is counted as an Out-of-order Packet.

Certain streaming protocol (such as TCP) require the packets to be in a certain order. Packets outside this are dropped by the streaming protocols even though there 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.

Out-of-order Packet count is based on the worst case streaming protocol. It allows for no reordering.

Packet loss does not affect the Out-of-order Packet count. Only packets that were not received in the order that they were transmitted.

Measurement units:
Packet count

See Also:
Stream
Test Sequence number

3.3.3 Duplicate Packet

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

Discussion:

Measurement units:
  Packet count

See Also:
Stream
Test Sequence number

3.3.4 Stream

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

Discussion:
Streams are tracked by test sequence number or "unique signature field" (RFC 2889). Streams define how individual packet's 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 transverse 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 mulitcast stream can be forward to multiple destination interfaces.

Measurement units:
n/a

See Also:
Flow
MicroFlow [3]
Test sequence number

3.3.6 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 and

the traffic receiver checks the value upon receipt of the packet. 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 number of received packets, the traffic receiver may also report number of in-sequence packets, number of out-sequence packets, number of duplicate packets, and number of reordered packets.

The recommended algorithm to use to change the sequence number on sequential packets is an incrementing value.

Measurement units:
n/a

See Also:
Stream

3.4 Vectors

A vector is a group of packets all containing a specific DSCP or IP precedence value. Vectors are expressed as a pair of numbers. The first is being the particular diff-serv value. The second is the metric expressed as a rate, loss percentage, delay, or jitter.

3.4.1 Intended Vector

Definition:

A vector describing the rate at which packets having a specific code-point (or IP precedence) that an external source attempts to transmit to a DUT/SUT.

Discussion:
Intended loads across the different code-point classes determine the metrics associated with a specific code-point traffic class.

Measurement Units:
N-octets packets per second

See Also:
Offered Vector
Expected Forwarding Vector
Expected Loss Vector
Expected Sequence Vector

Expected Delay Vector
Expected Jitter Vector
Forwarding Vector

Loss Vector

3.4.2 Offered Vector

Definition:
A vector describing the measured rate at which packets having a specific DSCP or IP precedence value are offered to the DUT/SUT.

Discussion:
Offered loads across the different code-point classes, constituting a code-point set, determine the metrics associated with a specific code-point traffic class.

Measurement Units:
N-octets packets per second

See Also:
Expected Forwarding Vector
Expected Loss Vector
Expected Sequence Vector
Expected Delay Vector
Expected Jitter Vector
Forwarding Vector
Codepoint Set

3.4.3 Expected Vectors

3.4.3.1 Expected Forwarding Vector

Definition:
A vector describing the expected output rate of packets having a specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:
The DUT is configured in a certain way in order that service differentiation occurs for a particular DSCP or IP precedence value when a specific traffic mix consisting of multiple DSCPs or IP precedence values are applied. This term attempts to capture the expected forwarding behavior when subjected to a certain offered vectors.

The actual algorithm or mechanism the DUT uses to achieve service differentiation is not important in describing the expected forwarding vector.

Measurement units:
  N-octet packets per second

See Also:
Intended Vector
Offered Vector
Output Vectors
Expected Loss Vector
Expected Sequence Vector
Expected Delay Vector
Expected Jitter Vector

3.4.3.2 Expected Loss Vector

Definition:
A vector describing the percentage of packets, having a specific DSCP or IP precedence value, that should not be forwarded. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:
The DUT is configured in a certain way in order that service differentiation occurs for a particular DSCP or IP precedence value when a specific traffic mix consisting of multiple DSCPs or IP precedence values are applied. This term attempts to capture the expected forwarding behavior when subjected to a certain offered vectors.

The actual algorithm or mechanism the DUT uses to achieve service differentiation is not important in describing the expected loss vector.

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

See Also:
Intended Vector
Offered Vector
Expected Forwarding Vector
Expected Sequence Vector
Expected Delay Vector
Expected Jitter Vector

3.2.3.3 Expected Sequence Vector

Definition:

A vector describing the expected in-sequence packets having a specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:

The DUT is configured in a certain way in order that service differentiation occurs for a particular DSCP or IP precedence value when a specific traffic mix consisting of multiple DSCPs or IP precedence values are applied. This term attempts to capture the expected forwarding behavior when subjected to a certain offered vectors.

The actual algorithm or mechanism the DUT uses to achieve service differentiation is not important in describing the expected sequence vector.

Measurement Units:

N-octet packets per second

See Also:

Intended Vector
Offered Vector
Output Vectors
Expected Loss Vector
Expected Forwarding Vector
Expected Delay Vector
Expected Jitter Vector

3.4.3.4 Expected Instantaneous Delay Vector

Definition:

A vector describing the expected delay for packets having a specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:

The DUT is configured in a certain way in order that service differentiation occurs for a particular DSCP or IP precedence value when a specific traffic mix consisting of multiple DSCPs or IP precedence values are applied. This term attempts to capture the expected forwarding behavior when subjected to a certain offered vectors.

The actual algorithm or mechanism the DUT uses to achieve service differentiation is not important in describing the expected delay vector.

Measurement units:

Seconds.

See Also:

Intended Vector
Offered Vector
Output Vectors
Expected Loss Vector
Expected Sequence Vector
Expected Forwarding Vector
Expected Jitter Vector

3.4.3.5 Expected Average Delay Vector

Definition:
A vector describing the expected average delay for packets having a specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:
The DUT is configured in a certain way in order that service differentiation occurs for a particular DSCP or IP precedence value when a specific traffic mix consisting of multiple DSCPs or IP precedence values are applied. This term attempts to capture the expected forwarding behavior when subjected to a certain offered vectors.

The actual algorithm or mechanism the DUT uses to achieve service differentiation is not important in describing the expected average delay vector.

Measurement units:
Seconds.

See Also:
Intended Vector
Offered Vector
Output Vectors
Expected Loss Vector
Expected Sequence Vector
Expected Forwarding Vector
Expected Jitter Vector

3.4.3.6 Expected Maximum Delay Vector

Definition:
A vector describing the expected maximum delay for packets having a specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:
The DUT is configured in a certain way in order that service differentiation occurs for a particular DSCP or IP precedence

value when a specific traffic mix consisting of multiple DSCPs or IP precedence values are applied. This term attempts to capture the expected forwarding behavior when subjected to a certain offered vectors.

The actual algorithm or mechanism the DUT uses to achieve service differentiation is not important in describing the expected maximum delay vector.

Measurement units:
Seconds.

See Also:
Intended Vector
Offered Vector
Output Vectors
Expected Loss Vector
Expected Sequence Vector
Expected Forwarding Vector
Expected Jitter Vector

3.4.3.7 Expected Minimum Delay Vector

Definition:
A vector describing the expected minimum delay for packets having a specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:
The DUT is configured in a certain way in order that service differentiation occurs for a particular DSCP or IP precedence value when a specific traffic mix consisting of multiple DSCPs or IP precedence values are applied. This term attempts to capture the expected forwarding behavior when subjected to a certain offered vectors.

The actual algorithm or mechanism the DUT uses to achieve service differentiation is not important in describing the expected minimum delay vector.

Measurement units:
Seconds.

See Also:
Intended Vector
Offered Vector
Output Vectors

Expected Loss Vector
Expected Sequence Vector
Expected Forwarding Vector

Expected Jitter Vector

3.2.3.8 Expected Instantaneous Jitter Vector

Definition:
A vector describing the expected jitter between two consecutive packets' arrival times having a specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

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

The 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 delay and i is the test sequence number. Packets lost are not counted in the measurement.

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:
Seconds

See Also:
Delay
Jitter
Offered Vector
Output Vectors
Expected Average Jitter Vector
Expected Peak-to-peak Jitter Vector
Stream

3.2.3.9 Expected Average Jitter Vector

Definition:
A vector describing the expected jitter in packet arrival times for packets having specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:
Average Jitter Vector is the average of all the Instantaneous

Jitter Vectors measured during the test duration for the same
DSCP or IP precedence value.

Measurement units:
Seconds

See Also:
Intended Vector
Offered Vector
Output Vectors
Expected Instantaneous Jitter Vector
Expected Peak-to-peak Jitter Vector

3.2.3.10 Expected Peak-to-peak Jitter Vector

Definition:
A vector describing the expected maximum variation in the delay of packet arrival times for packets having specific DSCP or IP precedence value. The value is dependent on the set of offered vectors and configuration of the DUT.

Discussion:
Peak-to-peak Jitter Vector is the maximum delay minus the minimum 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.

Measurement units:
Seconds

See Also:
Intended Vector
Offered Vector
Output Vectors
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 all packets containing a specific DSCP or IP precedence value that a device can be observed to successfully forward to the correct destination interface in response to an offered vector.

Discussion:
Forwarding Vector is expressed as pair of numbers. Both the specific DSCP (or IP precedence) value AND the packets per second value combine to make a vector.

The Forwarding Vector represents packet rate based on its specific DSCP (or IP precedence) value. It is not necessarily based on a stream or flow. The Forwarding Vector may be expressed as per port of the DUT/SUT. However, it must be consistent with the Expected Forwarding Vector.

Forwarding Vector is a per-hop measurement. The DUT/SUT may change the specific DSCP (or IP precedence) value for a multiple-hop measurement.

Measurement units:
N-octet packets per second

See Also:
Intended Vector
Offered Vector
Expected Vectors
Loss Vector
Sequence Vector
Delay Vectors

3.4.4.2 Loss Vector

Definition:
The percentage of packets containing specific DSCP or IP precedence value that a DUT/SUT did not transmit to the correct destination interface in response to an offered vector.

Discussion:
Loss Vector is expressed as pair of numbers. Both the specific DSCP (or IP precedence) value AND the percentage value combine to make a vector.

The Loss Vector represents percentage based on a specific DSCP or IP precedence value. It is not necessarily based on a stream or flow. The Loss Vector may be expressed as per port of the DUT/SUT. However, it must be consistent with the Expected Loss Vector

Loss Vector is a per-hop measurement. The DUT/SUT may change the specific DSCP or IP precedence value for a multiple-hop measurement.

Measurement Units:
Percentage of offered packets that are not forwarded.

See Also:
Intended Vector
Offered Vector
Expected Vectors
Forwarding Vector
Sequence Vector
Delay Vectors

3.4.4.3 Sequence Vector

Definition:
The number of packets per second for all packets containing a specific DSCP or IP precedence value that a device can be observed to transmit in sequence to the correct destination interface in response to an offered vector.

Discussion:
Sequence Vector is expressed as pair of numbers. Both the specific DSCP (or IP precedence) value AND the packets per second value combine to make a vector.

The Sequence Vector represents packet rate based on its specific DSCP or IP precedence value. It is not necessarily based on a stream or flow. The Sequence Vector may be expressed as per port of the DUT/SUT. However, it must be consistent with the Expected Sequence Vector.

Sequence Vector is a per-hop measurement. The DUT/SUT may change the specific DSCP or IP precedence value for a multiple-hop measurement.

Measurement Units:
N-octet packets per second

Issues:

See Also:
In-sequence Packet
Intended Vector
Offered Vector

Expected Vectors
Loss Vector
Forwarding Vector

Delay Vectors

3.4.4.4 Instantaneous Delay Vector

Definition:
The delay for a packet containing specific DSCP or IP precedence value that a device can be observed to successfully transmit to the correct destination interface in response to an offered vector.

Discussion:
Instantaneous Delay vector is expressed as pair of numbers. Both the specific DSCP (or IP precedence) value AND delay value combine to make a vector.

The Instantaneous Delay Vector represents delay on its specific DSCP or IP precedence value. It is not necessarily based on a stream or flow. The Delay vector may be expressed as per port of the DUT/SUT. However, it must be consistent with the Expected Delay vectors.

Instantaneous Delay Vector is a per-hop measurement. The DUT/SUT may change the specific DSCP or IP precedence value for a multiple-hop measurement.

Instantaneous Delay vector can be obtained at any offered load. Recommend at or below the channel capacity in the absence of congestion. For congested delay, run the offered load above the channel capacity.

Forwarding Vector may contain several Instantaneous Delay Vectors. For every packet received in a Forwarding Vector, there is a corresponding Instantaneous Delay Vector.

Measurement Units:
Seconds

See Also:
Delay
Intended Vector
Offered Vector
Expected Delay Vectors
Average Delay Vector
Maximum Delay Vector
Minimum Delay Vector

3.4.4.5 Average Delay Vector

Definition:

The average delay for packets containing specific DSCP or IP precedence value that a device can be observed to successfully transmit to the correct destination interface in response to an offered vector.

Discussion:

Average Delay vector is expressed as pair of numbers. Both the specific DSCP (or IP precedence) value AND delay value combine to make a vector.

The Delay Vector represents delay on its specific DSCP or IP precedence value. It is not necessarily based on a stream or flow. The Delay vector may be expressed as per port of the DUT/SUT. However, it must be consistent with the Expected Delay vector.

The Average Delay Vector is computed by averaging all the Instantaneous Delay Vectors for a given vector.

Average Delay Vector is a per-hop measurement. The DUT/SUT may change the specific DSCP or IP precedence value for a multiple-hop measurement.

Average Delay vector can be obtained at any offered load. Recommend at or below the channel capacity in the absence of congestion. For congested delay, run the offered load above the channel capacity.

Measurement Units:

Seconds

See Also:

Delay
Intended Vector
Offered Vector
Expected Delay Vectors
Instantaneous Delay Vector
Maximum Delay Vector
Minimum Delay Vector

3.4.4.6 Maximum Delay Vector

Definition:

The maximum delay from all packets containing specific DSCP or IP precedence value that a device can be observed to successfully transmit to the correct destination interface in response to an offered vector.

Discussion:

Maximum Delay vector is expressed as pair of numbers. Both the specific DSCP (or IP precedence) value AND delay value combine to make a vector.

The Maximum Delay Vector represents delay on its specific DSCP or IP precedence value. It is not necessarily based on a stream or flow. The Maximum Delay vector may be expressed as per port of the DUT/SUT. However, it must be consistent with the Expected Delay vector.

Maximum Delay Vector is based upon the maximum Instantaneous Delay Vector of all packets in a Forwarding Vector.

Maximum Delay Vector is a per-hop measurement. The DUT/SUT may change the specific DSCP or IP precedence value for a multiple-hop measurement.

Measurement Units:
Seconds

See Also:
Delay
Intended Vector
Offered Vector
Expected Delay Vectors
Instantaneous Delay Vector
Forwarding Vector
Average Delay Vector
Minimum Delay Vector

3.4.4.7 Minimum Delay Vector

Definition:
The minimum delay from all packets containing specific DSCP or IP precedence value that a device can be observed to successfully transmit to the correct destination interface in response to an offered vector.

Discussion:
Delay vector is expressed as pair of numbers. Both the specific DSCP (or IP precedence) value AND delay value combine to make a vector.

The Minimum Delay Vector represents delay on its specific DSCP or IP precedence value. It is not necessarily based on a stream or flow. The Minimum Delay vector may be expressed as per port of the DUT/SUT. However, it must be consistent with the Expected Delay vector.

Minimum Delay Vector is based upon the minimum Instantaneous Delay Vector of all packets in a Forwarding Vector.

Minimum Delay Vector is a per-hop measurement. The DUT/SUT may change the specific DSCP or IP precedence value for a multiple-hop measurement.

Minimum Delay vector can be obtained at any offered load. Recommend at or below the channel capacity in the absence of congestion. For congested delay, run the offered load above the channel capacity.

Measurement Units:
Seconds

See Also:
Delay
Intended Vector
Offered Vector
Expected Delay Vectors
Instantaneous Delay Vector
Forwarding Vector
Average Delay Vector
Maximum Delay Vector

3.4.4.8 Instantaneous Jitter Vector

Definition:
The jitter for two consecutive packets containing specific DSCP or IP precedence value that a device can be observed to successfully transmit to the correct destination interface in response to an offered vector.

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

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

The 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 delay and i is the test sequence number. Packets lost are not counted in the measurement.

Instantaneous Jitter Vector is a per-hop measurement. The

DUT/SUT may change the specific DSCP or IP precedence value for a multiple-hop measurement.

Forwarding Vector may contain several Instantaneous Jitter Vectors. For n packets received in a Forwarding Vector, there are exactly (n-1) Instantaneous Jitter Vectors.

Measurement units:
Seconds

See Also:
Delay
Jitter
Forwarding Vector
Stream
Expected Vectors
Average Jitter Vector
Peak-to-peak Jitter Vector

3.4.4.9 Average Jitter Vector

Definition:
The average jitter for packets containing specific DSCP or IP precedence value that a device can be observed to successfully transmit to the correct destination interface in response to an offered vector.

Discussion:
Average Jitter Vector is the average of all the Instantaneous Jitter Vectors of the same DSCP or IP precedence value, measured during the test duration.

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

Average Jitter vector is a per-hop measurement. The DUT/SUT may change the specific DSCP or IP precedence value for a multiple-hop measurement.

Measurement units:
Seconds

See Also:
Jitter
Forwarding Vector
Stream
Expected Vectors
Instantaneous Jitter Vector
Peak-to-peak Jitter Vector

3.4.4.10 Peak-to-peak Jitter Vector

Definition:
The maximum possible variation in the delay for packets containing specific DSCP or IP precedence value that a device can be observed to successfully transmit to the correct destination interface in response to an offered vector.

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

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

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:
Seconds

See Also:
Jitter
Forwarding Vector
Stream
Expected Vectors
Average Jitter Vector
Peak-to-peak 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. References

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

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

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

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

[5] V. Jacobson, K. Nichols, K. Poduri, _An Expedited Forwarding PHB_, RFC 2598, June 1999

[6] C. Demichelis, P. Chimento, _IP Packet Delay Variation Metric for IPPM_, draft-ietf-ippm-ipdv-10.txt

[7] H. Schulzrinne, GMD Fokus, S. Casner, R. Frederick, V. Jacobson, _RTP: A Transport Protocol for Real-Time Applications_, RFC 1889, January 1996

[8] A. Mankin, K. Ramakrishnan, _Gateway Congestion Control Survey_, RFC 1254, August 1991

6. Authors' Address

Jerry Perser
Spirent Communications
26750 Agoura Road
Calabasas, CA 91302
USA

Phone: + 1 818 676 2300
EMail: jerry.perser@spirentcom.com

David Newman
Network Test
31324 Via Colinas, Suite 113
Westlake Village, CA 91362
USA

Phone: + 1 818 889 0011, x10
EMail: dnewman@networktest.com

Sumit Khurana
Telcordia Technologies
445 South Street
Morristown, NJ 07960
USA

Phone: + 1 973 829 3170
EMail: sumit@research.telcordia.com

Shobha Erramilli
QNetworx Inc
1119 Campus Drive West
Morganville NJ 07751
USA

Phone:
EMail: shobha@qnetworx.com

Scott Poretsky
Avici Systems
101 Billerica Ave_Building #6
N. Billerica, MA 01862
USA

Phone: + 1 978 964 2287

EMail: sporetsky@avici.com

7. Full Copyright Statement