

Network Working Group
Request for Comments: 2889
Category: Informational

R. Mandeville
CQOS Inc.
J. Perser
Spirent Communications
August 2000

Benchmarking Methodology for LAN Switching Devices

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Table of Contents

1. Introduction	2
2. Requirements	2
3. Test setup	2
4. Frame formats and sizes	3
5. Benchmarking Tests	3
5.1 Fully meshed throughput, frame loss and forwarding rates	4
5.2 Partially meshed one-to-many/many-to-one	7
5.3 Partially meshed multiple devices	10
5.4 Partially meshed unidirectional traffic	13
5.5 Congestion Control	16
5.6 Forward Pressure and Maximum Forwarding Rate	19
5.7 Address caching capacity	22
5.8 Address learning rate	25
5.9 Errored frames filtering.	27
5.10 Broadcast frame Forwarding and Latency	28
6. Security Considerations	30
7. References	30
8. Authors' Addresses	30
Appendix A: Formulas	31
Appendix B: Generating Offered Load	32
Full Copyright Statement	35

1. Introduction

This document is intended to provide methodology for the benchmarking of local area network (LAN) switching devices. It extends the methodology already defined for benchmarking network interconnecting devices in RFC 2544 [3] to switching devices.

This RFC primarily deals with devices which switch frames at the Medium Access Control (MAC) layer. It provides a methodology for benchmarking switching devices, forwarding performance, congestion control, latency, address handling and filtering. In addition to defining the tests, this document also describes specific formats for reporting the results of the tests.

A previous document, "Benchmarking Terminology for LAN Switching Devices" [2], defined many of the terms that are used in this document. The terminology document SHOULD be consulted before attempting to make use of this document.

2. Requirements

The following RFCs SHOULD be consulted before attempting to make use of this document: RFC 1242 [1], RFC 2285 [2], and RFC 2544 [3].

For the sake of clarity and continuity, this RFC adopts the template for benchmarking tests set out in Section 26 of RFC 2544.

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. Test setup

This document extends the general test setup described in section 6 of RFC 2544 [3] to the benchmarking of LAN switching devices. RFC 2544 [3] primarily describes non-meshed traffic where input and output interfaces are grouped in mutually exclusive sending and receiving pairs. In fully meshed traffic, each interface of a DUT/SUT is set up to both receive and transmit frames to all the other interfaces under test.

Prior to each test run, the DUT/SUT MUST learn the MAC addresses used in the test and the address learning SHOULD be verified. Addresses not learned will be forwarded as flooded frames and reduce the amount of correctly forwarded frames. The rate at which address learning frames are offered may have to be adjusted to be as low as 50 frames per second or even less, to guarantee successful learning. The DUT/SUT address aging time SHOULD be configured to be greater than

the period of the learning phase of the test plus the trial duration plus any configuration time required by the testing device. Addresses SHOULD NOT age out until the trial duration is completed. More than one learning trial may be needed for the association of the address to the port to occur.

If a DUT/SUT uses a hashing algorithm with address learning, the DUT/SUT may not learn the necessary addresses to perform the tests. The format of the MAC addresses MUST be adjustable so that the address mapping may be re-arranged to ensure that the DUT/SUT learns all the addresses.

4. Frame formats and sizes

The test frame format is defined in RFC 2544 section 8 [3] and MUST contain a unique signature field located in the UDP DATA area of the Test Frame (see Appendix C [3]). The purpose of the signature field is filter out frames that are not part of the offered load.

The signature field MUST be unique enough to identify the frames not originating from the DUT/SUT. The signature field SHOULD be located after byte 56 (collision window [4]) or at the end of the frame. The length, contents and method of detection is not defined in this memo.

The signature field MAY have a unique identifier per port. This would filter out misforwarded frames. It is possible for a DUT/SUT to strip off the MAC layer, send it through its switching matrix, and transmit it out with the correct destination MAC address but the wrong payload.

For frame sizes, refer to RFC 2544, section 9 [3].

There are three possible frame formats for layer 2 Ethernet switches: standard MAC Ethernet frames, standard MAC Ethernet frames with vendor-specific tags added to them, and IEEE 802.3ac frames tagged to accommodate 802.1p&Q. The two types of tagged frames may exceed the standard maximum length frame of 1518 bytes, and may not be accepted by the interface controllers of some DUT/SUTs. It is recommended to check the compatibility of the DUT/SUT with tagged frames before testing.

Devices switching tagged frames of over 1518 bytes will have a different maximum forwarding rate than untagged frames.

5. Benchmarking Tests

The following tests offer objectives, procedures, and reporting formats for benchmarking LAN switching devices.

5.1 Fully meshed throughput, frame loss and forwarding rates

5.1.1 Objective

To determine the throughput, frame loss and forwarding rates of DUT/SUTs offered fully meshed traffic as defined in RFC 2285 [2].

5.1.2 Setup Parameters

When offering full meshed traffic, the following parameters MUST be defined. Each parameter is configured with the following considerations.

Frame Size - Recommended frame sizes are 64, 128, 256, 512, 1024, 1280 and 1518 bytes, per RFC 2544 section 9 [3]. The four CRC bytes are included in the frame size specified.

Interframe Gap (IFG) - The IFG between frames inside a burst MUST be at the minimum specified by the standard (9.6 us for 10Mbps Ethernet, 960 ns for 100Mbps Ethernet, and 96 ns for 1 Gbps Ethernet) of the medium being tested.

Duplex mode - Half duplex or full duplex.

ILoad - Intended Load per port is expressed in a percentage of the medium's maximum theoretical load, regardless of traffic orientation or duplex mode. Certain test configurations will theoretically over-subscribe the DUT/SUT.

In half duplex, an ILoad over 50% will over-subscribe the DUT/SUT.

Burst Size - The burst size defines the number of frames sent back-to-back at the minimum legal IFG [4] before pausing transmission to receive frames. Burst sizes SHOULD vary between 1 and 930 frames. A burst size of 1 will simulate constant load [1].

Addresses per port - Represents the number of addresses which are being tested for each port. Number of addresses SHOULD be a binary exponential (i.e. 1, 2, 4, 8, 16, 32, 64, 128, 256, ...). Recommended value is 1.

Trial Duration - The recommended Trial Duration is 30 seconds. Trial duration SHOULD be adjustable between 1 and 300 seconds.

5.1.3 Procedure

All ports on the tester MUST transmit test frames either in a Frame Based or Time Based mode (Appendix B). All ports SHOULD start transmitting their frames within 1% of the trial duration. For a trial duration of 30 seconds, all ports SHOULD have started transmitting frames within 300 milliseconds of each other.

Each port in the test MUST send test frames to all other ports in a round robin type fashion. The sequence of addresses MUST NOT change when congestion control is applied. The following table shows how each port in a test MUST transmit test frames to all other ports in the test. In this example, there are six ports with 1 address per port:

Source Port	Destination Ports (in order of transmission)					
Port #1	2	3	4	5	6	2...
Port #2	3	4	5	6	1	3...
Port #3	4	5	6	1	2	4...
Port #4	5	6	1	2	3	5...
Port #5	6	1	2	3	4	6...
Port #6	1	2	3	4	5	1...

As shown in the table, there is an equal distribution of destination addresses for each transmit opportunity. This keeps the test balanced so that one destination port is not overloaded by the test algorithm and all ports are equally and fully loaded throughout the test. Not following this algorithm exactly will produce inconsistent results.

For tests using multiple addresses per port, the actual port destinations are the same as described above and the actual source/destination address pairs SHOULD be chosen randomly to exercise the DUT/SUT's ability to perform address lookups.

For every address, learning frames MUST be sent to the DUT/SUT to allow the DUT/SUT update its address tables properly.

5.1.4 Measurements

Each port should receive the same number of test frames that it transmitted. Each receiving port MUST categorize, then count the frames into one of two groups:

- 1.) Received Frames: received frames MUST have the correct destination MAC address and SHOULD match a signature field.
- 2.) Flood count [2].

Any frame originating from the DUT/SUT (spanning tree, SNMP, RIP, ...) MUST not be counted as a received frame. Frames originating from the DUT/SUT MAY be counted as flooded frames or not counted at all.

Frame loss rate of the DUT/SUT SHOULD be reported as defined in section 26.3 [3] with the following notes: Frame loss rate SHOULD be measured at the end of the trail duration. The term "rate", for this measurement only, does not imply the units in the fashion of "per second."

5.1.4.1 Throughput

Throughput measurement is defined in section 26.1 [3]. A search algorithm is employed to find the maximum Oload [2] with a zero Frame loss rate [1]. The algorithm MUST adjust Iload to find the throughput.

5.1.4.2 Forwarding Rate

Forwarding rate (FR) of the DUT/SUT SHOULD be reported as the number of test frames per second that the device is observed to successfully forward to the correct destination interface in response to a specified Oload. The Oload MUST also be cited.

Forwarding rate at maximum offered load (FRMOL) MUST be reported as the number of test frames per second that a device can successfully transmit to the correct destination interface in response to the MOL as defined in section 3.6 [2]. The MOL MUST also be cited.

Maximum forwarding rate (MFR) MUST be reported as the highest forwarding rate of a DUT/SUT taken from an iterative set of forwarding rate measurements. The iterative set of forwarding rate measurements are made by adjusting Iload. The Oload applied to the device MUST also be cited.

5.1.5 Reporting format

The results for these tests SHOULD be reported in the form of a graph. The x coordinate SHOULD be the frame size, the y coordinate SHOULD be the test results. There SHOULD be at least two lines on the graph, one plotting the theoretical and one plotting the test results.

To measure the DUT/SUT's ability to switch traffic while performing many different address lookups, the number of addresses per port MAY be increased in a series of tests.

5.2 Partially meshed one-to-many/many-to-one

5.2.1 Objective

To determine the throughput when transmitting from/to multiple ports and to/from one port. As with the fully meshed throughput test, this test is a measure of the capability of the DUT to switch frames without frame loss. Results of this test can be used to determine the ability of the DUT to utilize an Ethernet port when switching traffic from multiple Ethernet ports.

5.2.2 Setup Parameters

When offering bursty meshed traffic, the following parameters MUST be defined. Each parameter is configured with the following considerations.

Frame Size - Recommended frame sizes are 64, 128, 256, 512, 1024, 1280 and 1518 bytes, per RFC 2544 section 9 [3]. The four CRC bytes are included in the frame size specified.

Traffic Direction - Traffic can be generated in one direction, the reverse direction, or both directions.

Interframe Gap (IFG) - The IFG between frames inside a burst MUST be at the minimum specified by the standard (9.6 us for 10Mbps Ethernet, 960 ns for 100Mbps Ethernet, and 96 ns for 1 Gbps Ethernet) of the medium being tested.

Duplex mode - Half duplex or full duplex.

ILoad - Intended Load per port is expressed in a percentage of the medium's maximum theoretical load, regardless of traffic orientation or duplex mode. Certain test configurations will theoretically over-subscribe the DUT/SUT.

In half duplex bidirectional traffic, an ILoad over 50% will over-subscribe the DUT/SUT.

Burst Size - The burst size defines the number of frames sent back-to-back at the minimum legal IFG [4] before pausing transmission to receive frames. Burst sizes SHOULD vary between 1 and 930 frames. A burst size of 1 will simulate constant load [1].

Addresses per port - Represents the number of addresses which are being tested for each port. Number of addresses SHOULD be a binary exponential (i.e. 1, 2, 4, 8, 16, 32, 64, 128, 256, ...). Recommended value is 1.


Trial Duration - The recommended Trial Duration is 30 seconds. Trial duration SHOULD be adjustable between 1 and 300 seconds.

5.2.3 Procedure

All ports on the tester MUST transmit test frames either in a Frame Based or Time Based mode (Appendix B). Depending upon traffic direction, some or all of the ports will be transmitting. All ports SHOULD start transmitting their frames within 1% of the trial duration. For a trial duration of 30 seconds, all ports SHOULD have started transmitting frames within 300 milliseconds of each other.

Test frames transmitted from the Many Ports MUST be destined to the One port. Test frames transmitted from the One Port MUST be destined to the Many ports in a round robin type fashion. See section 5.1.3 for a description of the round robin fashion.

For tests using multiple addresses per port, the actual port destinations are the same as described above and the actual source/destination address pairs SHOULD be chosen randomly to exercise the DUT/SUT's ability to perform address lookups.



For every address, the testing device MUST send learning frames to allow the DUT/SUT to update its address tables properly.

5.2.4 Measurements

Each receiving port MUST categorize, then count the frames into one of two groups:

- 1.) Received Frames: received frames MUST have the correct destination MAC address and SHOULD match a signature field.
- 2.) Flood count [2].

Any frame originating from the DUT/SUT MUST not be counted as a received frame. Frames originating from the DUT/SUT MAY be counted as flooded frames or not counted at all.

Forwarding rate (FR) of the DUT/SUT SHOULD be reported as the number of test frames per second that the device is observed to successfully transmit to the correct destination interface in response to a specified Oload. The Oload MUST also be cited.

Forwarding rate at maximum offered load (FRMOL) MUST be reported as the number of test frames per second that a device can successfully transmit to the correct destination interface in response to the MOL as defined in section 3.6 [2]. The MOL MUST also be cited.

Maximum forwarding rate (MFR) MUST be reported as the highest forwarding rate of a DUT/SUT taken from an iterative set of forwarding rate measurements. The iterative set of forwarding rate measurements are made by adjusting Iload. The Oload applied to the device MUST also be cited.

5.2.5 Reporting Format

The results for these tests SHOULD be reported in the form of a graph. The x coordinate SHOULD be the frame size, the y coordinate SHOULD be the test results. There SHOULD be at least two lines on the graph, one plotting the theoretical and one plotting the test results.

To measure the DUT/SUT's ability to switch traffic while performing many different address lookups, the number of addresses per port MAY be increased in a series of tests.

5.3 Partially meshed multiple devices

5.3.1 Objective

To determine the throughput, frame loss and forwarding rates of two switching devices equipped with multiple ports and one high speed backbone uplink (Gigabit Ethernet, ATM, SONET).

5.3.2 Setup Parameters

When offering bursty partially meshed traffic, the following parameters **MUST** be defined. Each variable is configured with the following considerations.

Frame Size - Recommended frame sizes are 64, 128, 256, 512, 1024, 1280 and 1518 bytes, per RFC 2544 section 9 [3]. The four CRC bytes are included in the frame size specified.

Interframe Gap (IFG) - The IFG between frames inside a burst **MUST** be at the minimum specified by the standard (9.6 us for 10Mbps Ethernet, 960 ns for 100Mbps Ethernet, and 96 ns for 1 Gbps Ethernet) of the medium being tested.

Duplex mode - Half duplex or full duplex.

ILoad - Intended Load per port is expressed in a percentage of the medium's maximum theoretical load, regardless of traffic orientation or duplex mode. Certain test configurations will theoretically over-subscribe the DUT/SUT.

In half duplex, an ILoad over 50% will over-subscribe the DUT/SUT.

Burst Size - The burst size defines the number of frames sent back-to-back at the minimum legal IFG [4] before pausing transmission to receive frames. Burst sizes **SHOULD** vary between 1 and 930 frames. A burst size of 1 will simulate constant load [1].

Addresses per port - Represents the number of addresses which are being tested for each port. Number of addresses **SHOULD** be a binary exponential (i.e. 1, 2, 4, 8, 16, 32, 64, 128, 256, ...). Recommended value is 1.

Trial Duration - The recommended Trial Duration is 30 seconds. Trial duration **SHOULD** be adjustable between 1 and 300 seconds.

Local Traffic - A Boolean value of ON or OFF. The frame sequence algorithm MAY be altered to remove local traffic. With local traffic ON, the algorithm is exactly the same as a fully meshed throughput. With local traffic OFF, the port sends frames to all other ports on the other side of the backbone uplink in a round robin type fashion.

5.3.3 Procedure

All ports on the tester MUST transmit test frames either in a Frame Based or Time Based mode (Appendix B). All ports SHOULD start transmitting their frames within 1% of the trial duration. For a trial duration of 30 seconds, all ports SHOULD have started transmitting frames with 300 milliseconds of each other.

Each port in the test MUST send test frames to all other ports in a round robin type fashion as defined in section 5.1.3. Local traffic MAY be removed from the round robin list in order to send the entire load across the backbone uplink.

For tests using multiple addresses per port, the actual port destinations are the same as described above and the actual source/destination address pairs SHOULD be chosen randomly to exercise the DUT/SUT's ability to perform address lookups.

For every address, the testing device MUST send learning frames to allow the DUT/SUT to update its address tables properly.

To measure the DUT/SUT's ability to switch traffic while performing many different address lookups, the number of addresses per port MAY be increased in a series of tests.

5.3.4 Measurements

Each receiving port MUST categorize, then count the frames into one of two groups:

- 1.) Received frames MUST have the correct destination MAC address and SHOULD match a signature field.
- 2.) Flood count [2].

Any frame originating from the DUT/SUT MUST not be counted as a received frame. Frames originating from the DUT/SUT MAY be counted as flooded frames or not counted at all.

Frame loss rate of the DUT/SUT SHOULD be reported as defined in section 26.3 [3] with the following notes: Frame loss rate SHOULD be measured at the end of the trial duration. The term "rate", for this measurement only, does not imply the units in the fashion of "per second."

5.3.4.1 Throughput

Throughput measurement is defined in section 26.1 [3]. A search algorithm is employed to find the maximum Oload [2] with a zero Frame loss rate [1]. The algorithm MUST adjust Iload to find the throughput.

5.3.4.2 Forwarding rate

Forwarding rate (FR) of the DUT/SUT SHOULD be reported as the number of test frames per second that the device is observed to successfully forward to the correct destination interface in response to a specified Oload. The Oload MUST also be cited.

Forwarding rate at maximum offered load (FRMOL) MUST be reported as the number of test frames per second that a device can successfully transmit to the correct destination interface in response to the MOL as defined in section 3.6 [2]. The MOL MUST also be cited.

Maximum forwarding rate (MFR) MUST be reported as the highest forwarding rate of a DUT/SUT taken from an iterative set of forwarding rate measurements. The iterative set of forwarding rate measurements are made by adjusting Iload. The Oload applied to the device MUST also be cited.

5.3.5 Reporting format

The results for these tests SHOULD be reported in the form of a graph. The x coordinate SHOULD be the frame size, the y coordinate SHOULD be the test results. There SHOULD be at least two lines on the graph, one plotting the theoretical and one plotting the test results.

To measure the DUT/SUT's ability to switch traffic while performing many different address lookups, the number of addresses per port MAY be increased in a series of tests.

5.4 Partially meshed unidirectional traffic

5.4.1 Objective

To determine the throughput of the DUT/SUT when presented multiple streams of unidirectional traffic with half of the ports on the DUT/SUT are transmitting frames destined to the other half of the ports.

5.4.2 Setup Parameters

The following parameters MUST be defined. Each variable is configured with the following considerations.

Frame Size - Recommended frame sizes are 64, 128, 256, 512, 1024, 1280 and 1518 bytes, per RFC 2544 section 9 [3]. The four CRC bytes are included in the frame size specified.

Interframe Gap (IFG) - The IFG between frames inside a burst MUST be at the minimum specified by the standard (9.6 us for 10Mbps Ethernet, 960 ns for 100Mbps Ethernet, and 96 ns for 1 Gbps Ethernet) of the medium being tested.

Duplex mode - Half duplex or full duplex.

ILoad - Intended Load per port is expressed in a percentage of the medium's maximum theoretical load, regardless of traffic orientation or duplex mode. Certain test configurations will theoretically over-subscribe the DUT/SUT.

ILoad will not over-subscribe the DUT/SUT in this test.

Burst Size - The burst size defines the number of frames sent back-to-back at the minimum legal IFG [4] before pausing transmission to receive frames. Burst sizes SHOULD vary between 1 and 930 frames. A burst size of 1 will simulate constant load [1].

Addresses per port - Represents the number of addresses which are being tested for each port. Number of addresses SHOULD be a binary exponential (i.e. 1, 2, 4, 8, 16, 32, 64, 128, 256, ...). Recommended value is 1.

Trial Duration - The recommended Trial Duration is 30 seconds. Trial duration SHOULD be adjustable between 1 and 300 seconds.

5.4.3 Procedure

Ports do not send and receive test frames simultaneously. As a consequence, there should be no collisions unless the DUT is misforwarding frames, generating flooded or Spanning-Tree frames or is enabling some flow control mechanism. Ports used for this test are either transmitting or receiving, but not both. Those ports which are transmitting send test frames destined to addresses corresponding to each of the ports receiving. This creates a unidirectional mesh of traffic.

All ports on the tester MUST transmit test frames either in a Frame Based or Time Based mode (Appendix B). All ports SHOULD start transmitting their frames within 1% of the trial duration. For a trial duration of 30 seconds, all ports SHOULD have started transmitting frames with 300 milliseconds of each other.

Each transmitting port in the test MUST send frames to all receiving ports in a round robin type fashion. The sequence of addresses MUST NOT change when congestion control is applied. The following table shows how each port in a test MUST transmit test frames to all other ports in the test. In this 8 port example, port 1 through 4 are transmitting and ports 5 through 8 are receiving; each with 1 address per port:

Source Port, then Destination Ports (in order of transmission)

Port #1	5	6	7	8	5	6...
Port #2	6	7	8	5	6	7...
Port #3	7	8	5	6	7	8...
Port #4	8	5	6	7	8	5...

As shown in the table, there is an equal distribution of destination addresses for each transmit opportunity. This keeps the test balanced so that one destination port is not overloaded by the test algorithm and all receiving ports are equally and fully loaded throughout the test. Not following this algorithm exactly will product inconsistent results.

For tests using multiple addresses per port, the actual port destinations are the same as described above and the actual source/destination address pairs SHOULD be chosen randomly to exercise the DUT/SUT's ability to perform address lookups.

For every address, the testing device MUST send learning frames to allow the DUT/SUT to load its address tables properly. The address table's aging time SHOULD be set sufficiently longer than

the learning time and trial duration time combined. If the address table ages out during the test, the results will show a lower performing DUT/SUT.

To measure the DUT/SUT's ability to switch traffic while performing many different address lookups, the number of addresses per port MAY be increased in a series of tests.

5.4.4 Measurements

Each receiving port MUST categorize, then count the frames into one of two groups:

- 1.) Received Frames: received frames MUST have the correct destination MAC address and SHOULD match a signature field.
- 2.) Flood count [2].

Any frame originating from the DUT/SUT MUST not be counted as a received frame. Frames originating from the DUT/SUT MAY be counted as flooded frames or not counted at all.

Frame loss rate of the DUT/SUT SHOULD be reported as defined in section 26.3 [3] with the following notes: Frame loss rate SHOULD be measured at the end of the trial duration. The term "rate", for this measurement only, does not imply the units in the fashion of "per second."

5.4.4.1 Throughput

Throughput measurement is defined in section 26.1 [3]. A search algorithm is employed to find the maximum Oload [2] with a zero Frame loss rate [1]. The algorithm MUST adjust Iload to find the throughput.

5.4.4.2 Forwarding rate

Forwarding rate (FR) of the DUT/SUT SHOULD be reported as the number of test frames per second that the device is observed to successfully forward to the correct destination interface in response to a specified Oload. The Oload MUST also be cited.

Forwarding rate at maximum offered load (FRMOL) MUST be reported as the number of test frames per second that a device can successfully transmit to the correct destination interface in response to the MOL as defined in section 3.6 [2]. The MOL MUST also be cited.

Maximum forwarding rate (MFR) MUST be reported as the highest forwarding rate of a DUT/SUT taken from an iterative set of forwarding rate measurements. The iterative set of forwarding rate measurements are made by adjusting Iload. The Oload applied to the device MUST also be cited.

5.4.5 Reporting format

The results for these tests SHOULD be reported in the form of a graph. The x coordinate SHOULD be the frame size, the y coordinate SHOULD be the test results. There SHOULD be at least two lines on the graph, one plotting the theoretical and one plotting the test results.

To measure the DUT/SUT's ability to switch traffic while performing many different address lookups, the number of addresses per port MAY be increased in a series of tests.

5.5 Congestion Control

5.5.1 Objective

To determine how a DUT handles congestion. Does the device implement congestion control and does congestion on one port affect an uncongested port. This procedure determines if Head of Line Blocking and/or Backpressure are present.

5.5.2 Setup Parameters

The following parameters MUST be defined. Each variable is configured with the following considerations.

Frame Size - Recommended frame sizes are 64, 128, 256, 512, 1024, 1280 and 1518 bytes, per RFC 2544 section 9 [3]. The four CRC bytes are included in the frame size specified.

Interframe Gap (IFG) - The IFG between frames inside a burst MUST be at the minimum specified by the standard (9.6 us for 10Mbps Ethernet, 960 ns for 100Mbps Ethernet, and 96 ns for 1 Gbps Ethernet) of the medium being tested.


Duplex mode - Half duplex or full duplex.

Addresses per port - Represents the number of addresses which are being tested for each port. Number of addresses SHOULD be a binary exponential (i.e. 1, 2, 4, 8, 16, 32, 64, 128, 256, ...). Recommended value is 1.

Trial Duration - The recommended Trial Duration is 30 seconds.
 Trial duration SHOULD be adjustable between 1 and 300 seconds.

5.5.3 Procedure

This test MUST consist of a multiple of four ports with the same MOL. Four ports are REQUIRED and MAY be expanded to fully utilize the DUT/SUT in increments of four. Each group of four will contain a test block with two of the ports as source transmitters and two of the ports as receivers. The diagram below depicts the flow of traffic between the switch ports:



Both source transmitters MUST transmit the exact number of test frames. The first source MUST transmit test frames at the MOL with the destination address of the two receive ports in an alternating order. The first test frame to the uncongested receive port, second test frame to the congested receive port, then repeat. The second source transmitter MUST transmit test frames at the MOL only to the congested receive port.

Both receive ports SHOULD distinguish between test frames originating from the source ports and frames originating from the DUT/SUT. Only test frames from the source ports SHOULD be counted.

The uncongested receive port should be receiving at a rate of half the MOL. The number of test frames received on the uncongested port SHOULD be 50% of the test frames transmitted by the first source transmitter. The congested receive port should be receiving at the MOL. The number of test frames received on the congested port should be between 100% and 150% of the test frames transmitted by one source transmitter.

Test frames destined to uncongested ports in a switch device should not be dropped due to other ports being congested, even if the source is sending to both the congested and uncongested ports.

5.5.4 Measurements

Any frame received which does not have the correct destination address MUST not be counted as a received frame and SHOULD be counted as part of a flood count.

Any frame originating from the DUT/SUT MUST not be counted as a received frame. Frames originating from the DUT/SUT MAY be counted as flooded frames or not counted at all.

Frame loss rate of the DUT/SUT's congested and uncongested ports MUST be reported as defined in section 26.3 [3] with the following notes: Frame loss rate SHOULD be measured at the end of the trial duration. The term "rate", for this measurement only, does not imply the units in the fashion of "per second."

Offered Load to the DUT/SUT MUST be reported as the number of test frames per second that the DUT/SUT observed to accept. This may be different than the MOL.

Forwarding rate (FR) of the DUT/SUT's congested and uncongested ports MUST be reported as the number of test frames per second that the device is observed to successfully transmit to the correct destination interface in response to a specified offered load. The offered load MUST also be cited.

5.5.5 Reporting format

This test MUST report the frame lost rate at the uncongested port, the forwarding rate (at 50% offered load) at the uncongested port, and the frame lost rate at the congested port. This test MAY report the frame counts transmitted and frame counts received by the DUT/SUT.

5.5.5.1 HOLB

If there is frame loss at the uncongested port, "Head of Line" blocking is present. The DUT cannot forward the amount of traffic to the congested port and as a result it is also losing frames destined to the uncongested port.

5.5.5.2 Back Pressure

If there is no frame loss on the congested port, then backpressure is present. It should be noted that this test expects the overall load to the congested port to be greater than 100%. Therefore if the load is greater than 100% and no frame loss is detected, then the DUT must be implementing a flow control mechanism. The type of flow control mechanism used is beyond the scope of this memo.

It should be noted that some DUTs may not be able to handle the 100% load presented at the input port. In this case, there may be frame loss reported at the uncongested port which is due to the load at the input port rather than the congested port's load.

If the uncongested frame loss is reported as zero, but the maximum forwarding rate is less than 7440 (for 10Mbps Ethernet), then this may be an indication of congestion control being enforced by the DUT. In this case, the congestion control is affecting the throughput of the uncongested port.

If no congestion control is detected, the expected percentage frame loss for the congested port is 33% at 150% overload. It is receiving 100% load from 1 port, and 50% from another, and can only get 100% possible throughput, therefore having a frame loss rate of 33% (150%-50%/150%).

5.6 Forward Pressure and Maximum Forwarding Rate

5.6.1 Objective

The Forward Pressure test overloads a DUT/SUT port and measures the output for forward pressure [2]. If the DUT/SUT transmits frames with an interframe gap less than 96 bits (section 4.2.3.2.2 [4]), then forward pressure is detected.

The objective of the Maximum Forwarding Rate test is to measure the peak value of the Forwarding Rate when the Offered Load is varied between the throughput [1] and the Maximum Offered Load [2].

5.6.2 Setup Parameters

The following parameters MUST be defined. Each variable is configured with the following considerations.

Frame Size - Recommended frame sizes are 64, 128, 256, 512, 1024, 1280 and 1518 bytes, per RFC 2544 section 9 [3]. The four CRC bytes are included in the frame size specified.

Duplex mode - Half duplex or full duplex.

Trial Duration - The recommended Trial Duration is 30 seconds.
Trial duration SHOULD be adjustable between 1 and 300 seconds.

Step Size - The minimum incremental resolution that the Iload will be incremented in frames per second. The smaller the step size, the more accurate the measurement and the more iterations required. As the Iload approaches the MOL, the minimum step size will increase because of gap resolution on the testing device.

5.6.3 Procedure

5.6.3.1 Maximum forwarding rate

If the Throughput [1] and the MOL [2] are the same, then MFR [2] is equal to the MOL [2].

This test MUST at a minimum be performed in a two-port configuration as described below. Learning frames MUST be sent to allow the DUT/SUT to update its address tables properly.

Test frames are transmitted to the first port (port 1) of the DUT/SUT at the Iload. The FR [2] on the second port (port 2) of the DUT/SUT is measured. The Iload is incremented for each Step Size to find the MFR. The algorithm for the test is as follows:

```
CONSTANT
  MOL = ... frames/sec; {Maximum Offered Load}
VARIABLE
  MFR := 0 frames/sec; {Maximum Forwarding Rate}
  ILOAD := starting throughput in frames/sec; {offered load}
  STEP := ... frames/sec; {Step Size}
BEGIN
  ILOAD := ILOAD - STEP;
  DO
  BEGIN
    ILOAD := ILOAD + STEP
    IF (ILOAD > MOL) THEN
    BEGIN
      ILOAD := MOL
    END
    AddressLearning; {Port 2 broadcasts with its source address}
    Transmit(ILOAD); {Port 1 sends frames to Port 2 at Offered load}
    IF (Port 2 Forwarding Rate > MFR) THEN
    BEGIN
      MFR := Port 2 Forwarding Rate; {A higher value than before}
    END
```

```
END
WHILE (ILOAD < MOL); {ILOAD has reached the MOL value}
DONE
```

5.6.3.2 Minimum Interframe Gap

The Minimum Interframe gap test SHOULD, at a minimum, be performed in a two-port configuration as described below. Learning frames MUST be sent to allow the DUT/SUT to update its address tables properly.

Test frames SHOULD be transmitted to the first port (port 1) of the DUT/SUT with an interframe gap of 88 bits. This will apply forward pressure to the DUT/SUT and overload it at a rate of one byte per frame. The test frames MUST be constructed with a source address of port 1 and a destination address of port 2.

The FR on the second port (port 2) of the DUT/SUT is measured. The measured Forwarding Rate should not exceed the medium's maximum theoretical utilization (MOL).

5.6.4 Measurements

Port 2 MUST categorize, then count the frames into one of two groups:

- 1.) Received Frames: received frames MUST have the correct destination MAC address and SHOULD match a signature field.
- 2.) Flood count [2].

Any frame originating from the DUT/SUT MUST not be counted as a received frame. Frames originating from the DUT/SUT MAY be counted as flooded frames or not counted at all.

5.6.5 Reporting format

MFR MUST be reported as the highest forwarding rate of a DUT/SUT taken from an iterative set of forwarding rate measurements. The Iload applied to the device MUST also be cited.

Forwarding rate (FR) of the DUT/SUT SHOULD be reported as the number of frames per second that the device is observed to successfully transmit to the correct destination interface in response to a specified Oload. The Iload MUST be cited and the Oload MAY be recorded.

If the FR exceeds the MOL during the Minimum Interframe gap test, this MUST be highlighted with the expression "Forward Pressure detected".

5.7 Address Caching Capacity

5.7.1 Objective

To determine the address caching capacity of a LAN switching device as defined in RFC 2285, section 3.8.1 [2].

5.7.2 Setup Parameters

The following parameters MUST be defined. Each variable is configured with the following considerations.

Age Time - The maximum time that a DUT/SUT will keep a learned address in its forwarding table.

Addresses Learning Rate - The rate at which new addresses are offered to the DUT/SUT to be learned. The rate at which address learning frames are offered may have to be adjusted to be as low as 50 frames per second or even less, to guarantee successful learning.

Initial Addresses - The initial number of addresses to start the test with. The number MUST be between 1 and the maximum number supported by the implementation.

5.7.3 Procedure

The aging time of the DUT/SUT MUST be known. The aging time MUST be longer than the time necessary to produce frames at the specified rate. If a low frame rate is used for the test, then it may be possible that sending a large amount of frames may actually take longer than the aging time.

This test MUST at a minimum be performed in a three-port configuration described below. The test MAY be expanded to fully utilized the DUT/SUT in increments of two or three ports. An increment of two would include an additional Learning port and Test port. An increment of three would include an additional Learning port, Test port, and Monitoring port.

The Learning port (Lport) transmits learning frames to the DUT/SUT with varying source addresses and a fixed destination address corresponding to the address of the device connected to the Test port (Tport) of the DUT/SUT. By receiving frames with varying source addresses, the DUT/SUT should learn these new addresses. The source addresses MAY be in sequential order.

The Test port (Tport) of the DUT/SUT acts as the receiving port for the learning frames. Test frames will be transmitted back to the addresses learned on the Learning port. The algorithm for this is explained below.

The Monitoring port (Mport) on the DUT/SUT acts as a monitoring port to listen for flooded or mis-forwarded frames. If the test spans multiple broadcast domains (VLANs), each broadcast domain REQUIRES a Monitoring port.

It is highly recommended that SNMP, Spanning Tree, and any other frames originating from the DUT/SUT be disabled when running this test. If such protocols cannot be turned off, the flood count MUST be modified only to count test frame originating from Lport and MUST NOT count frames originating from the DUT/SUT.

The algorithm for the test is as follows:

```

CONSTANT
    AGE = ...;    {value greater than DUT aging time}
    MAX = ...;    {maximum address support by implementation}
VARIABLE
    LOW  := 0;    {Highest passed value}
    HIGH := MAX;  {Lowest failed value}
    N    := ...;  {user specified initial starting point}
BEGIN
    DO
        BEGIN
            PAUSE(AGE);    {Age out any learned addresses}
            AddressLearning(TPort); {broadcast a frame with its source
                                   Address and broadcast destination}
            AddressLearning(LPort); {N frames with varying source addresses
                                   to Test Port}
            Transmit(TPort); {N frames with varying destination addresses
                              corresponding to Learning Port}
            IF (MPort receive frame != 0) OR
                (LPort receive frames < TPort transmit) THEN
                BEGIN {Address Table of DUT/SUT was full}
                    HIGH := N;
                END
            ELSE
                BEGIN {Address Table of DUT/SUT was NOT full}
                    LOW := N;
                END
            N := LOW + (HIGH - LOW)/2;
        END WHILE (HIGH - LOW >= 2);
    END {Value of N equals number of addresses supported by DUT/SUT}

```

Using a binary search approach, the test targets the exact number of addresses supported per port with consistent test iterations. Due to the aging time of DUT/SUT address tables, each iteration may take some time during the waiting period for the addresses to clear. If possible, configure the DUT/SUT for a low value for the aging time.

Once the high and low values of N meet, then the threshold of address handling has been found.

5.7.4 Measurements

Whether the offered addresses per port was successful forwarded without flooding.

5.7.5 Reporting format

After the test is run, results for each iteration SHOULD be displayed in a table to include:

The number of addresses used for each test iteration (varied).

The intended load used for each test iteration (fixed).

Number of test frames that were offered to Tport of the DUT/SUT. This SHOULD match the number of addresses used for the test iteration. Test frames are the frames sent with varying destination addresses to confirm that the DUT/SUT has learned all of the addresses for each test iteration.

The flood count on Tport during the test portion of each test. If the number is non-zero, this is an indication of the DUT/SUT flooding a frame in which the destination address is not in the address table.

The number of frames correctly forwarded to test Lport during the test portion of the test. Received frames MUST have the correct destination MAC address and SHOULD match a signature field. For a passing test iteration, this number should be equal to the number of frames transmitted by Tport.

The flood count on Lport during the test portion of each test. If the number is non-zero, this is an indication of the DUT/SUT flooding a frame in which the destination address is not in the address table.

The flood count on Mport. If the value is not zero, then this indicates that for that test iteration, the DUT/SUT could not determine the proper destination port for that many frames. In other words, the DUT/SUT flooded the frame to all ports since its address table was full.

5.8 Address Learning Rate

5.8.1 Objective

To determine the rate of address learning of a LAN switching device.

5.8.2 Setup Parameters

The following parameters MUST be defined. Each variable is configured with the following considerations.

Age Time - The maximum time that a DUT/SUT will keep a learned address in its forwarding table.

Initial Addresses Learning Rate - The starting rate at which new addresses are offered to the DUT/SUT to be learned.

Number of Addresses - The number of addresses that the DUT/SUT must learn. The number MUST be between 1 and the maximum number supported by the implementation. It is recommended no to exceed the address caching capacity found in section 5.9

5.8.3 Procedure

The aging time of the DUT/SUT MUST be known. The aging time MUST be longer than the time necessary to produce frames at the specified rate. If a low frame rate is used for the test, then it may be possible that sending a large amount of frames may actually take longer than the aging time.

This test MUST at a minimum be performed in a three-port configuration in section 5.9.3. The test MAY be expanded to fully utilized the DUT/SUT in increments of two or three ports. An increment of two would include an additional Learning port and Test port. An increment of three would include an additional Learning port, Test port, and Monitoring port.

An algorithm similar to the one used to determine address caching capacity can be used to determine the address learning rate. This test iterates the rate at which address learning frames are offered

by the test device connected to the DUT/SUT. It is recommended to set the number of addresses offered to the DUT/SUT in this test to the maximum caching capacity.

The address learning rate might be determined for different numbers of addresses but in each test run, the number MUST remain constant and SHOULD be equal to or less than the maximum address caching capacity.

5.8.4 Measurements

Whether the offered addresses per port were successful forwarded without flooding at the offered learning rate.

5.8.5 Reporting format

After the test is run, results for each iteration SHOULD be displayed in a table:

The number of addresses used for each test iteration (fixed).

The intended load used for each test iteration (varied).

Number of test frames that were transmitted by Tport. This SHOULD match the number of addresses used for the test iteration. Test frames are the frames sent with varying destination addresses to confirm that the DUT/SUT has learned all of the addresses for each test iteration.

The flood count on Tport during the test portion of each test. If the number is non-zero, this is an indication of the DUT/SUT flooding a frame in which the destination address is not in the address table.

The number of frames correctly forwarded to test Lport during the test portion of the test. Received frames MUST have the correct destination MAC address and SHOULD match a signature field. For a passing test iteration, this number should be equal to the number of frames transmitted by Tport.

The flood count on Lport during the test portion of each test. If the number is non-zero, this is an indication of the DUT/SUT flooding a frame in which the destination address is not in the address table.

The flood count on Mport. If the value is not zero, then this indicates that for that test iteration, the DUT/SUT could not determine the proper destination port for that many frames. In other words, the DUT/SUT flooded the frame to all ports since its address table was full.

5.9 Errored frames filtering

5.9.1 Objective

The objective of the Errored frames filtering test is to determine the behavior of the DUT under error or abnormal frame conditions. The results of the test indicate if the DUT/SUT filters the errors, or simply propagates the errored frames along to the destination.

5.9.2 Setup Parameters

The following parameters MUST be defined. Each variable is configured with the following considerations.

ILoad - Intended Load per port is expressed in a percentage of the medium's maximum theoretical load possible. The actual transmitted frame per second is dependent upon half duplex or full duplex operation. The test SHOULD be run multiple times with a different load per port in each case.

Trial Duration - The recommended Trial Duration is 30 seconds. Trial duration SHOULD be adjustable between 1 and 300 seconds.

5.9.3 Procedure

Each of the illegal frames for Ethernet MUST be checked:

Oversize - The DUT/SUT MAY filter frames larger than 1518 bytes from being propagated through the DUT/SUT section 4.2.4.2.1 [4]. Oversized frames transmitted to the DUT/SUT should not be forwarded. DUT/SUT supporting tagged Frames MAY forward frames up to and including 1522 bytes long (section 4.2.4.2.1 [5]).

Undersize - The DUT/SUT MUST filter frames less than 64 bytes from being propagated through the DUT/SUT (section 4.2.4.2.2 [4]). Undersized frames (or collision fragments) received by the DUT/SUT must not be forwarded.

CRC Errors - The DUT/SUT MUST filter frames that fail the Frame Check Sequence Validation (section 4.2.4.1.2 [4]) from being propagated through the DUT/SUT. Frames with an invalid CRC transmitted to the DUT/SUT should not be forwarded.

Dribble Bit Errors - The DUT/SUT MUST correct and forward frames containing dribbling bits. Frames transmitted to the DUT/SUT that do not end in an octet boundary but contain a valid frame check sequence MUST be accepted by the DUT/SUT (section 4.2.4.2.1 [4]) and forwarded to the correct receive port with the frame ending in an octet boundary (section 3.4 [4]).

Alignment Errors - The DUT/SUT MUST filter frames that fail the Frame Check Sequence Validation AND do not end in an octet boundary. This is a combination of a CRC error and a Dribble Bit error. When both errors are occurring in the same frame, the DUT/SUT MUST determine the CRC error takes precedence and filters the frame (section 4.2.4.1.2 [4]) from being propagated.

5.9.5 Reporting format

For each of the error conditions in section 5.6.3, a "pass" or "fail" MUST be reported. Actual frame counts MAY be reported for diagnostic purposes.

5.10 Broadcast frame Forwarding and Latency

5.10.1 Objective

The objective of the Broadcast Frame Forwarding and Latency Test is to determine the throughput and latency of the DUT when forwarding broadcast traffic. The ability to forward broadcast frames will depend upon a specific function built into the device for that purpose. It is therefore necessary to determine the ability of DUT/SUT to handle broadcast frames, since there may be many different ways of implementing such a function.

5.10.2 Setup Parameters

The following parameters MUST be defined. Each variable is configured with the following considerations.

Frame Size - Recommended frame sizes are 64, 128, 256, 512, 1024, 1280 and 1518 bytes, per RFC 2544 section 9 [3]. The four CRC bytes are included in the frame size specified.

Duplex mode - Half duplex or full duplex.

ILoad - Intended Load per port is expressed in a percentage of the medium's maximum theoretical load, regardless of traffic orientation or duplex mode. Certain test configurations will theoretically over-subscribe the DUT/SUT.

ILoad will not over-subscribe the DUT/SUT in this test.

Trial Duration - The recommended Trial Duration is 30 seconds.
Trial duration SHOULD be adjustable between 1 and 300 seconds.

5.10.3 Procedure

For this test, there are two parts to be run.

Broadcast Frame Throughput - This portion of the test uses a single source test port to transmit test frames with a broadcast address using the frame specified in RFC 2544 [3]. Selected receive ports then measure the forwarding rate and Frame loss rate.

Broadcast Frame Latency - This test uses the same setup as the Broadcast Frame throughput, but instead of a large stream of test frames being sent, only one test frame is sent and the latency to each of the receive ports are measured in seconds.

5.10.4 Measurements

Frame loss rate of the DUT/SUT SHOULD be reported as defined in section 26.3 [3] with the following notes: Frame loss rate SHOULD be measured at the end of the trial duration. The term "rate", for this measurement only, does not imply the units in the fashion of "per second."

Forwarding rate (FR) of the DUT/SUT SHOULD be reported as the number of test frames per second that the device is observed to successfully forward to the correct destination interface in response to a specified Oload. The Oload MUST also be cited.

5.10.5 Reporting format

The results for these tests SHOULD be reported in the form of a graph. The x coordinate SHOULD be the frame size, the y coordinate SHOULD be the test results. There SHOULD be at least two lines on the graph, one plotting the theoretical and one plotting the test results.

To measure the DUT/SUT's ability to switch traffic while performing many different address lookups, the number of addresses per port MAY be increased in a series of tests.

6. Security Considerations

As this document is solely for the purpose of providing metric methodology and describes neither a protocol nor a protocol's implementation, there are no security considerations associated with this document.

7. 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] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.
- [4] ANSI/IEEE, "CSMA/CD Access Method and Physical Layer Specifications," ISO/IEC 8802-3, ISBN 0-7381-0330-6, 1998.
- [5] IEEE Draft, "Frame Extensions for Virtual Bridged Local Area Networks (VLAN) Tagging on 802.3 Networks", 802.3ac/D3.1, July 1998.

8. Authors' Addresses

Robert Mandeville
CQOS Inc.
21 Technology
Irvine, CA 92618
USA

Phone: +1 (949) 400-4444
EMail: bob@cqos.com

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

Phone: + 1 818 676 2300
EMail: jerry_perser@netcomsystems.com

Appendix A: Formulas

A.1 Calculating the InterBurst Gap

IBG is defined in RFC 2285 [2] as the interval between two bursts. To achieve a desired load, the following Input Parameter need to be defined:

- LENGTH - Frame size in bytes including the CRC.
- LOAD - The intended load in percent. Range is 0 to 100.
- BURST - The number of frames in the burst (integer value).
- SPEED - media's speed in bits/sec
 - Ethernet is 10,000,000 bits/sec
 - Fast Ethernet is 100,000,000 bits/sec
 - Gigabit Ethernet is 1,000,000,000 bits/sec
- IFG - A constant 96 bits for the minimum interframe gap.

The IBG (in seconds) can be calculated:

$$\text{IBG} = \frac{[(100/\text{LOAD} - 1) * \text{BURST} * (\text{IFG} + 64 + 8 * \text{LENGTH})] + \text{IFG}}{\text{SPEED}}$$

A.2 Calculating the Number of Bursts for the Trial Duration

The number of bursts for the trial duration is rounded up to the nearest integer number. The follow Input Parameter need to be defined:

- LENGTH - Frame size in bytes including the CRC.
- BURST - The number of frames in the burst (integer value).
- SPEED - media's speed in bits/sec
 - Ethernet is 10,000,000 bits/sec
 - Fast Ethernet is 100,000,000 bits/sec
 - Gigabit Ethernet is 1,000,000,000 bits/sec
- IFG - A constant 96 bits for the minimum interframe gap.
- IBG - Found in the above formula
- DURATION - Trial duration in seconds.

An intermediate number of the Burst duration needs to be calculated first:

$$\text{TXTIME} = \frac{\text{DURATION}}{\text{SPEED}}$$

Number of Burst for the Trial Duration (rounded up):

$$\text{\#OFBURSTS} = \frac{\text{DURATION}}{(\text{TXTIME} + \text{IBG})}$$

Example:

```

LENGTH   = 64  bytes per frame
LOAD      = 100 % offered load
BURST     = 24  frames per burst
SPEED     = 10  Mbits/sec (Ethernet)
DURATION  = 10  seconds test

```

```

IBG       = 1612.8 uS
TXTIME    = 1603.2 uS
\#OFBURSTS = 3110

```

Appendix B: Generating Offered Load

In testing, the traffic generator is configured with the Iload (Intended Load) and measures the Oload (Offered Load). If the DUT/SUT applies congestion control, then the Iload and the Oload are not the same value. The question arises, how to generate the Oload? This appendix will describe two different methods.

The unit of measurement for Oload is bits per second. The two methods described here will hold one unit constant and let the DUT/SUT vary the other unit. The traffic generator SHOULD specify which method it uses.

B.1 Frame Based Load

Frame Based Load holds the number of bits constant. The Trial Duration will vary based upon congestion control. Advantage is implementation is a simple state machine (or loop). The disadvantage is that Oload needs to be measured independently.

All ports on the traffic generator MUST transmit the exact number of test frames. The exact number of test frames is found by multiplying the Iload of the port by the Trial Duration. All ports MAY NOT transmit the same number of frames if their Iload is not the same. An example would be the Partially meshed many-to-one test.

All ports SHOULD start transmitting their frames within 1% of the trial duration. For a trial duration of 30 seconds, all ports SHOULD have started transmitting frames within 300 milliseconds of each other.

The reported Oload SHOULD be the average during the Trial Duration. If the traffic generator continues to transmit after the Trial Duration due to congestion control, Oload MAY be averaged over the entire transmit time. Oload for the DUT/SUT MUST be the aggregate of all the Oloads per port. Oload per port MAY be reported.

B.2 Time Based Load

Time based load holds the Trial Duration constant, while allowing the number of octets transmitted to vary. Advantages are an accurate Trial Duration and integrated Oload measurement. Disadvantage is that the starting and stopping of the traffic generator MUST be more accurate.

All ports on the traffic generator are configured to transmit the Iload for a finite amount of time. Each port MUST count the number of octets successfully transmitted.

The start and stop is initiated at a layer defined by the test parameters. The layer can be the MAC layer, IP layer, or some other point in the protocol stack. The traffic generator MUST complete its layer specific transmit process when the stop time is reached (i.e. no fragments, finish the frame).

All ports MUST start transmitting their frames within 1% of the trial duration. For a trial duration of 30 seconds, all ports SHOULD have started transmitting frames within 300 milliseconds of each other.

All ports SHOULD stop transmitting frames after the specified trial duration within 0.01% of the trial duration. Each port's stop time MUST be reference to its start time. This trial duration error controls the accuracy of the Oload measurement and SHOULD be reported with the Oload measurement.

Each port is allowed an offset error of 0.1% and a trial duration error of 0.01%.

Oload is found by taking the number of octets successfully transmitted and dividing by the trial duration. Oload for the DUT/SUT MUST be the aggregate of all the Oloads per port. Oload per port MAY be reported for diagnostic purposes.

Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

