Benchmarking Methodology Working Group Internet Draft - December 1995

Bob Mandeville ENL Ajay V. Shah Wandel & Goltermann Technologies, Inc

Benchmarking Methodology for Ethernet Switches

<<u>draft-ietf-bmwg-ethernet-switches-00.txt</u>>

Status of this Document

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet- Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the Internet-Drafts Shadow Directories on ds.internic.net (US East Coast), nic.nordu.net (Europe), ftp.isi.edu (US West Coast), or munnari.oz.au (Pacific Rim).

Distribution of this document is unlimited. Please send comments to bmwg@harvard.edu or to the editors.

Abstract

This initial draft sets out to define a methodology designed specifically for Ethernet switches. Although the roots of these devices are clearly to be found in bridging, switches have matured enough in the last few years to deserve some special attention. The test methodology described here concentrates on four areas: throughput, address handling, latency and behavior in abnormal conditions. This draft will attempt to make clear why switch specific tests are needed in each of these four areas and will define a number of tests in each area. In addition to defining the tests this document also describes specific formats for reporting the results of the tests. <u>Appendix A</u> lists the tests and conditions that we believe should be included for specific cases and gives additional information about testing practices. <u>Appendix B</u> is a reference listing of maximum frame rates to be used with specific frame sizes on Ethernet.

<u>1</u>. Introduction

Vendors often engage in "specsmanship" in an attempt to give their products a better position in the marketplace. This often involves "smoke & mirrors" to confuse the potential users of the products. This document and follow up memos attempt to define a specific set of tests that vendors can use to measure and report the performance characteristics of their ethernet switch. The results of these tests will provide the user comparable data from different vendors with which to evaluate these devices.

A previous document, "Benchmarking Terminology for Network Interconnect Devices" (<u>RFC 1242</u>), 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. Real world

Please refer to the draft on "Benchmarking Methodology for Network Interconnect Devices".

<u>3</u>. Tests to be run

Please refer to the draft on "Benchmarking Methodology for Network Interconnect Devices".

<u>4</u>. Evaluating the results

Please refer to the draft on "Benchmarking Methodology for Network Interconnect Devices".

5. Requirements

In this document, the words that are used to define the significance of each particular requirement are capitalized. These words are:

* "MUST"

This word or the adjective "REQUIRED" means that the item is an absolute requirement of the specification.

* "SHOULD"

This word or the adjective "RECOMMENDED" means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood and the case carefully weighed before choosing a different course.

* "MAY"

This word or the adjective "OPTIONAL" means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

An implementation is not compliant if it fails to satisfy one or more of the MUST requirements for the test it implements. An implementation that satisfies all the MUST and all the SHOULD requirements for the test is said to be "unconditionally compliant"; one that satisfies all the MUST requirements but not all the SHOULD requirements for the test is said to be "conditionally compliant".

<u>6</u>. Device set up

The device MUST be in a stable state (i.e.: the device SHOULD have completed its initialization process) prior to the start of any tests. Before starting to perform the tests, the device to be tested MUST be configured following the instructions provided to the user. Specifically, it is expected that all of the supported features will be configured and enabled during this set up (See Appendix A). It is expected that all of the tests will be run without changing the configuration or setup of the device in any way other than that required to do the specific test. For example, it is not acceptable to change the size of frame handling buffers between tests of frame handling rates when testing the throughput of the device. It is necessary to modify the configuration when starting a test to determine the effect of filters on throughput, but the only change MUST be to enable the specific filter. The device set up SHOULD include the normally recommended configuration. The specific version of the software and the exact device configuration, including what device functions are disabled, used during the tests SHOULD be included as part of the report of the results.

7. Frame sizes

All of the described tests SHOULD be performed at a number of frame sizes. Specifically, the sizes SHOULD include the maximum and minimum legitimate sizes for Ethernet and enough sizes in between to be able to get a full characterization of the device performance.

In most cases it makes more sense to test the device with the minimum frame size for the media since this would stress the device to its limits and help characterize the per-frame processing overhead of the device. However, the latency of the device under test MUST be evaluated using a range of different frame sizes as highlighted in <u>section 7.1</u>.

7.1 Frame sizes to be used on Ethernet

64, 128, 256, 512, 1024, 1280, 1518

These sizes include the maximum and minimum frame sizes permitted by the Ethernet standard and a selection of sizes between these extremes with a finer granularity for the smaller frame sizes and higher frame rates.

8. Modifiers

It might be useful to know the device performance under a number of conditions; some of these conditions are noted below. It is expected that the reported results will include as many of these conditions as the test equipment is able to generate. The suite of tests SHOULD be first run without any modifying conditions and then repeated under each of the conditions separately. To preserve the ability to compare the

results of these tests it is necessary to let the device know of the various addresses configured on all of its ports. A phenomenon know as the "Learning Process" is required for this purpose. This process MUST send any frames that are required to generate all the addresses to be used in the test, thus not requiring the device to learn addresses every time the first instance of the test is run. All of the addresses SHOULD resolve to the same "next-hop" and it is expected that this will be the address of the receiving side of the test equipment. This learning process will have to be repeated at the beginning of each test.

8.1 Broadcast frames

Please refer to the draft on "Benchmarking Methodology for Network Interconnect Devices".

9. Multidirectional traffic

Normal network activity is not all in a single direction. To test the multidirectional performance of a device, the test series SHOULD be run with the same data rate being offered from all the directions. The sum of the data rates should not exceed the theoretical limit for Ethernet.

<u>10</u>. Single stream path

The full suite of tests SHOULD be run along with whatever modifier conditions that are relevant using a single input and output network port on the device. If the internal design of the device has multiple distinct pathways, for example, multiple interface cards each with multiple network ports, then it is not necessary that all possible types of pathways SHOULD be tested separately since these tests test the basic switching fabric of the device and is not geared towards testing the internal hardware neither is it geared towards testing if the interface card is properly in place in the device.

<u>11</u>. Multiple frame sizes

This document does not address the issue of testing the effects of a mixed frame size environment other than to suggest that if such tests are wanted then frames SHOULD be distributed between all of the listed sizes for Ethernet. The distribution MAY approximate the conditions on the network in which the device would be used.

<u>12</u>. Maximum frame rate

The maximum frame rate that should be used when testing LAN connections SHOULD be the listed theoretical maximum rate for the frame size on Ethernet. A list of maximum frame rates for LAN connections is included in <u>Appendix B</u>.

<u>13</u>. Bursty traffic

It is convenient to measure the device performance under steady state load but this is an unrealistic way to gage the functioning of a device since actual network traffic normally consists of bursts of frames. Some of the tests described below SHOULD be performed with both steady state traffic and with traffic consisting of repeated bursts of frames. The frames within a burst are transmitted with the minimum legitimate inter-frame gap.

<u>14</u>. Trial description

A particular test consists of multiple trials. Each trial returns one piece of information, for example the loss rate at a particular input frame rate. Each trial consists of a number of phases:

a) Send the "learning frames" to the switch port and wait 2 seconds to be sure that the learning has settled. The formats of the learning frame that should be used are shown in the Test Frame Formats document.

- b) Run the test trial.
- c) Wait for two sec for any residual frames to be received.
- d) Wait for at least five seconds for the device to restabilize.

<u>15</u>. Trial duration

The aim of these tests is to determine the rate continuously supportable by the device. The actual duration of the test trials must be a compromise between this aim and the duration of the benchmarking test suite. The duration of the test portion of each trial SHOULD be at least 10 seconds.

<u>16</u>. Benchmarking tests:

Note: The notation "type of data stream" refers to the above modifications to a frame stream with a constant inter-frame gap, for example, the addition of traffic filters to the configuration of the device under test.

<u>16.1</u> Switch Throughput Test

Objective: To determine the throughput of the switching device under test.

Typically Ethernet switches are equipped with a large number of **10 Mbit/s** ports. Smaller devices often have 8 or 12 ports while larger devices equipped with special internal buses are equipped with over a hundred dedicated 10 Mbit/s ports. Such switching devices are built to forward Ethernet frames from their source to their destination addresses by establishing a connection between appropriate incoming and outgoing ports of the switching device. Since the source and destination ports through which traffic is forwarded are not predictable, it seems most appropriate to test the throughput of Ethernet switching devices by creating multiple streams of traffic with all of the ports on the device both sending and receiving frames at the same time. It also seems appropriate to call this kind of throughput as switching throughput. Closer examination of the types of multiple streams of traffic that might be created to test the switching throughput of a device makes it obvious that choices have to be made in order to make the testing process as efficient and pertinent as possible. These choices pertain to:

- a) the pattern of traffic to be created between the ports under test,
- b) the load of the test traffic,
- c) the inter-frame gap and the interval between bursts of the test traffic,
- d) the frame sizes used for the test traffic and
- e) the address learning process.

We will now look at each of these in turn.

a) The pattern of traffic to be created between the switching ports under test.

Almost all Ethernet switching devices can assign a large number of MAC addresses to each one of their ports. Test frames can therefore contain a large number of source addresses and destination addresses. This makes the number of connections that a switch can be required to set up during a test very large. A switch with just six ports and eight MAC addresses assigned to each would have to establish 1,920 (5x8x8x6) separate connections in order to forward frames between all of the source and destination addresses. It is necessary to keep in mind that the order in which individual connections between source and destination addresses need to be set up by the device under test will have a direct effect on the way in which the load of the test traffic is distributed over a devices ports. Consequently, in order to keep tight control over the distribution of the load presented to the ports under test it is desirable to create a pattern of multi-directional traffic which makes it possible to guarantee the exact loads that will result from traffic being generated in a multi-directional traffic pattern.

Summary:

It is desirable to create a multi-directional traffic pattern to test the throughput of Ethernet switching devices to ensure that the device under test creates switched connections on a frame-by-frame basis between all of the learned source and destination addresses on all of the tested ports. The multi-directional pattern of traffic has to be designed so that traffic switched in and out of the ports under test equals whatever target load is set.

Procedure:

The multi-directional traffic pattern can be achieved on a frame-by-frame basis simply as follows:

The first frame sent to port 1 from the test instrument SHOULD carry a destination address assigned to port 2. The second frame sent to port 1 SHOULD carry a destination address assigned to port 3 and so on until

the highest port number is reached. This pattern will then be repeated until the end of the test. Traffic will be sent to all of the ports under test in similar fashion. In this way all of the ports are both sending and receiving frames for the duration of the test. The same number of MAC addresses SHOULD be learned by each port for any one test. Naturally each port SHOULD learn a different set of such addresses. If the device allows and the test concentrates on switching throughput and not address handling then eight MAC addresses per port appears to be a good number to work with. If eight addresses are listed as belonging to a given port then the MAC addresses assigned to frames sent to that port can be picked randomly from the list for the duration of the test. The same can hold for the destination addresses that belong to each of the ports under test. The frames generated for the switching throughput test repeatatively rotate through the source and destination ports in an ascending order but the source and destination addresses behind each port can be chosen randomly. For example, the first frame will be sent to port 1 and forwarded to port 2. It will carry any one of the source addresses assigned to port 1 in its source address field and any one of the destination addresses assigned to port 2 in its destination field.

b) The load of the test traffic.

Ethernet is a half-duplex technology. Consequently the maximum load of **10** Mbit/s that a port can handle includes both the frames transmitted and the frames received by that port. Since the multi-directional pattern of traffic described above requires ports to send and receive frames at the same time, the only way to fully load all ports is to send exactly half the maximum number of frames per second that a port can legally handle to each of the ports under test. The load on each port is to be calculated in frames per second. The maximum number of frames per second that a 10 Mbit/s port can handle is a function of frame length (see Appendix B). For the multi-directional traffic pattern if the frame size is 64 bytes, the minimum legal length for Ethernet frames, then the total number of frames transmitted by any one port MUST not exceed 14880/2 frames. Because each port transmits in a round-robin fashion to all of the other ports the maximum number of frames received by any port will be (14880/2/number of ports) x the number of ports or 7440 frames. At maximum load therefore each port under test will see 7440 64-byte frames in and 7440 64-byte frames out. For loads inferior to 100% the same formula applies and guarantees that the load is evenly distributed over all of the ports under test. For example to achieve a 80% load on all switching ports under test the test instrument MUST send 40% of the maximum legal load for the frame length used for the test to each of the ports. The switch test SHOULD be run at different loads. Experience shows that loads of 70%, 80%, 90% and 100% provide significant results.

Overload Condition

Loads exceeding 100% can readily be created by this testing methodology since at the full line rate on short frames for example, a port is said to be loaded at 100% when it is transmitting at 50% of the line rate or 7440 frames per second and receiving at 50% of the line rate. Since ports can transmit at more than 50% of the full line rate, it is relatively easy to create loads well in excess of the Ethernet legal limit. It is important to build such overloads into the test methodology since they can occur on switched Ethernet networks when many ports send frames to a single receiving port at high rates.

Overloads SHOULD be part of the test methodology since they allow the tester to discover if the device under test implements some sort of back-pressure mechanism and also allows the tester to appreciate the size and efficiency of the buffers of device under test.

Overloads where the ports under test are required to process frames in excess of the legal maximum defined for each frame length of 10%, 20% and 30% provide significant results.

c) The inter-frame gap and the interval between bursts of the test traffic.

Experience demonstrates that real LAN traffic is often bursty in nature that is frames are mostly transmitted in dense packs with pauses occurring between successive packs. This makes it particularly desirable to construct a test pattern with bursty traffic where the number of frames in each burst and the time between bursts can be treated as variables. It SHOULD be noted that because ports both send and receive frames during the switching throughput test the test instrument cannot send sustained bursts of frames with minimum inter-frame gap that last the entire duration of the test as would be the case if the traffic was uni-directional and frames were sent from ports whose only job it is to transmit frames to ports whose only job it is to receive frames.

Procedure:

Traffic sent to ports SHOULD send frames in bursts such that the inter-frame gap between frames within a burst is set to the legal minimum of 9.6 microseconds and the interval between bursts is set according to the target load. This test methodology makes frequent reference to the interval between bursts which will from here on in be called the inter-burst gap or IBG.

We will assume that bursts contain frames of a single length and that the number of frames per burst is constant for any one test. IBG is a function of the line speed, which for Ethernet is 10 Mbit/s, the length of the frames in the bursts, the number of frames in the bursts, the length of the preamble and the inter-frame gap. If the target load is known the IBG can be calculated as follows: mmax (mediamaximum) = maximum number of bits per second supported by the media which for Ethernet is 10 000 000 bits per second. The media transmits one bit every 1 / 10 000 000 seconds. We will call this one bit time. ifg (inter-frame gap) = 96 bits for frames within a burst or 9.6 microseconds or 96 bit times. pre (preamble) = 64 bits or the time required to transmit a 64 bit Ethernet preamble on the media which equals 6.4 microseconds or 64 bit times. frt (frame time) = the time required to transmit one Ethernet frame of a given length on the media. A 64-byte frame is made up of 64 x 8 bytes or 512 bits and therefore requires 51.2 microseconds or 512 bit times to be transmitted on the 10 000 000 bit per second Ethernet media. For simplicity let FRT = pre + ifg + frt expressed in bit times. mfr (maximum frame rate) = maximum number of frames per second supported by the media which is a function of mmax and FRT: mfr = mmax / FRT. mfr is expressed in seconds. tload (target load) = the target load expressed as a percentage of the maximum number of frames the media can handle for a given frame length. For simplicity let mfr x tload/2 be designated as the sendrate, that is the number of frames per second that the test instrument SHOULD send to each one of the ports under test. bsize (burst size) = the number of frames contained in a single burst Using this notation and knowing what the target load is we can calculate the inter-burst gap as: IBG = (1 - (sendrate x FRT)) / ((sendrate / bsize) - 1) Here are two examples to illustrate. Example 1. To calculate the IBG required to generate a target load of 100% on all ports using 64-byte frames with 20 frames per burst we have: mmax = 10 000 000tload = 100%. We know that:

ifg = 96 bit times pre = 64 bit times frt = 512 bit times which add up to an FRT of 672 bit times. Substituting these values into mfr = mmax / FRT we get 10 000 000 / 672 = 14 880 so that the sendrate of mfr x tload/2 = 14 880 x 100%/2 = 7440. Furthermore, $7440 \times FRT = 5000000$ bit times or 0.5 seconds. Since each burst contains 20 frames and 7440 frames per second are sent to each port the number of bursts per second will equal 7440 / 20 = 372 bursts and the number of intervals between bursts will be equal the number of bursts minus one or 371. So we have: IBG = (1 - 0.5) / (372 - 1) = 0.5 / 371 = 0.0013477 seconds or 1.3477 milliseconds. Checking we find that 372 bursts of 20 frame transmissions taking 67.2 microseconds each add up to 0.5 seconds while 371 inter-burst gaps of 1.3477 milliseconds account for the remaining 0.5 seconds. Example 2. To calculate the IBG for a target load of 80% with 512-byte frames we first find the sendrate = 940 and the FRT = 4256 bit times. Then IBG = $(1s - (940 \times 425.6)) / ((940/128) - 1) = 94.58$ milliseconds. Checking we have 940 frames transmitted with a transmission time of 425.6 equal to 0.4 seconds and 6.34375 inter-burst gaps of 94.58 milliseconds equal to 0.06 seconds. The examples above show that not all combinations of frame size, burst size and load make a perfect fit because they fill each second of transmission time with fractions of the desired IBG. Given frame size and desired load it is however possible to calculate the characteristic values of the burst sizes which allow a perfect fit. The characteristic values for burst sizes less than 1000 which allow a perfect fit of the calculated IBG into every second of transmission time for 64-byte frames and for loads of 50% to 130% by 10% steps are: 1, 2, 3, 4, 5, 6, 8, 12, 24, 31, 62, 93, 124, 186, 248, 372, 744, 930. **Recommendations:** The size of the bursts SHOULD range from a small number of frames, 10, to a large number of frames, 500, or more. The perfect fit values for burst size are: 1, 2, 3, 4, 5, 6, 8, 12, 24, 31, 62, 93, 124, 186, 248, 372, 744, 930 frames. Experience tends to show that tests run along the lines of the above description do not require a large number of ports on a device to be put to test. The reason for this is that the switch specific test tends to power out the buffering schemes and transmission and receive engines of the devices under test long before putting any strain of the internal high speed buses that are built into this family of devices. The pattern of traffic created in the switch

test obviously results in a large number of collisions being produced during the test since ports are sending and receiving frames in all directions simultaneously. It is absolutely necessary to use a test generator which implements the Ethernet back-off algorithm. Even so as ports contend for access to the media they will be forced to buffer frames which could not be transmitted because of collisions. This puts the buffering mechanisms of the devices under test to work in the same way they would on real networks although the test methodology described here allows the tester to create full line loads and lengthy bursts which would only occur on a real network in extreme instances.

A note on the the back-off algorithm:

We have already pointed out the necessity for the test instrument sending traffic to the ports of switching devices under test to implement the Ethernet back-off algorithm. It SHOULD ne noted in addition that this algorithm uses the formula 2(**n) -1 where n ranges from 1 to 10 to calculate the number of 51.2 periods the transmitter SHOULD hold off before attempting a retransmission. It would appear some that Ethernet controllers limit the range of n to values less than 10. This results in a more aggressive back-off algorithm and has a direct consequence on the number of collision battles a device will win. It is therefore desirable to ascertain the relative aggressivity of the test instruments and of the devices under test when performing the multi-directional traffic test since many collisions occur during this kind of test.

d) The frame sizes used for the test traffic SHOULD conform to the frame sizes defined in $\frac{\text{RFC } 1242}{2}$.

e) The address learning process SHOULD conform to the process described in <u>RFC 1242</u>.

Reporting format:

Since the test devices and the devices under test MAY implement different back-off algorithms, it is impossible to determine how many frames will actually be sent to the ports of the device under test in advance. This makes it necessary to count the number of frames actually sent to each of the ports under test. This count MUST NOT include the number of collisions that occurred during the test. The number of frames received by each of the ports MUST also be counted. Using these two counts it becomes possible to determine the total percentage of frames transmitted without loss as follows:

Values calculated in this way can be reported in a table with lines representing different loads and rows representing different burst sizes:

burst sizes **20 40 128 256 500**

loads 70% 80% 90% 100% 110% 120% 130%

<u>16.2</u> Behavior Test

Objective:

To determine the behavior of the switching device under abnormal frame conditions.

Summary:

It is imperative to determine how an Ethernet switch would behave when it encounters frames which are defined to be illegal according to the Ethernet specifications. Such frames could be:

Runts Jabbers Bad FCS

As a special case, it MAY be worthwhile to determine how the switch behaves when it has to handle broadcasts as well as various sized Bad FCS frames.

Procedure:

Any one port on the switch MAY be chosen as the transmitter of such illegal frames and any other port MAY be selected as the receiver. Traffic containing these illegal frames MUST be sent from the transmitter and the traffic received by the receiver MUST be noted. If the receiver receives over 90% the illegal frames, we can qualify the switch to be forwarding the illegal frames otherwise the switch MAY be deemed as blocking or filtering the illegal frames.

Reporting Format: The results generated would be whether the switch is forwarding these illegal frames or filtering them.

<u>16.3</u> Latency tests

Objective: To measure the time it takes a frame to go through a switching device.

Summary:

A large proportion of switching devices do not store frames sent to them before retransmiiting them onto the media. This significantly reduces the amount of time frames take to go through the device under test and makes it desirable to have a test procedure which makes it easy to recognize whether a switch stores frames before forwarding them onto the media or not. This is also important since devices which store frames will have significantly greater latencies as frame lengths increase than devices which do not. Since there is no standard way for switches to process broadcast frames it is also very desirable to measure latency on broadcast frames.

Procedure:

The measurement of latency requires test instruments to place a time-stamp on the frames sent to the device under test. The time stamp or tag must be placed at the head of a frame when the frame is sent to a port of the switching device under test. The test instrument must also be able to record the time at which the head of the frame containing the time tag is received inbound on the test instrument once it has been retransmitted by the device under test back to the test instrument. Latency is calculated by subtracting this recorded receive time from the recorded send time on the time tag. This simple delta is used to express the latency of the device under test. Latency should be measured for unicast frames of standard test lengths, for broadcast frames and should be measured with different external loads applied to the device under test.

Reporting format:

The latency measurements are to be reported in microseconds in a table.

<u>16.4</u> Address handling test

Objective:

To determine how many addresses can be assigned to the ports of an Ethernet switching device.

Summary:

It is necessary to verify how many MAC addresses can be assigned to each port of a switching device before the device begins to drop frames.

Procedure:

The test must be set up to isolate the address handling capacity of the devices under test from their switching performance. In order to achieve this a simple uni-directional pattern of traffic with parallel streams of traffic is created. Half of the ports under test transmit frames to the other half of the ports which receive them. The test procedure is in two steps. In step 1 a small number of addresses is assigned to each port and in step 2 a stream of frames is sent back-to-back to the sending ports of the test and the number of frames received on the receiving ports counted. Steps 1 and 2 should be repeated until the receive count indicates that the device has dropped frames. Reporting format:

Roh Mandavilla

The number of addresses assigned to each port in the step before the first frames were dropped in step 2 of the test is to be presented as the number of addresses the switching device under test supports.

<u>17</u>. Editor's Addresses

ENL Labs	Phone: +33 1 4761 3057
ENL	Fax: +33 1 4278 3671
Paris, France	Email: bob.mandeville@eunet.fr
Ajay V. Shah	
Wandel & Goltermann, Inc	Phone: +1 919 941 5740
<u>P</u> . O. Box 13585	Fax: +1 919 941 5751
Research Triangle Park, NC 27709	Email: shah@wg.com
	Compuserve: 102755,712

Appendix A: Testing Considerations

A.1 Scope Of This Appendix

This appendix discusses certain issues in the ethernet switch testing methodology where experience or judgment may play a role in the tests selected to be run or in the approach to constructing the test with a particular device. As such, this <u>appendix M</u>UST not be read as an amendment to the methodology described in the body of this document but as a guide to testing practice.

1. Typical testing practice has been to wait till the device under test stabilizes after "boot-up", otherwise the switch may be generating frames like routing updates etc which may interfere with the test frames. It is also necessary to turn off the SNMP feature (if the switch has one) to make sure that the management frames don't interfere with the test frames.

<u>2</u>. The device under test MUST be configured before starting any test and the configuration MUST remain the same through out the testing phase.

3. Architectural considerations may need to be considered. For example, first perform the tests with the stream going between ports on the same interface card and the repeat the tests with the stream going into a port on one interface card and out of a port on a second interface card. There will almost always be a best case and worst case configuration for a given device under test architecture.

<u>4</u>. Testing done using traffic streams consisting of the smallest allowable frame size for the media has shown to be the most stressful on the devices under test.

Appendix B: Maximum frame rates reference Ethernet Size Ethernet (bytes) (pps) 64 14880 128 8445 256 4528 512 2349 768 1586 1024 1197 1280 961 1518 812 Ethernet size Preamble 64 bits Frame 8 x N bits Gap 96 bits Benchmarking Methodology Working Group Internet Draft - December 1995 Bob Mandeville ENL Ajay V. Shah Wandel & Goltermann Technologies, Inc B. Mandeville, A. V. Shah [Page 13] Expires in six months