

Understanding and Testing Gigabit Ethernet Networks

by Aaron Deragon

TABLE OF CONTENTS:

<u>1.0 What Is Ethernet?</u>	<u>02</u>
<u>2.0 Ethernet Frame Structure</u>	<u>02</u>
<u>3.0 Emerging Applications for Gigabit Ethernet</u>	<u>03</u>
<u>4.0 BERT and Ethernet</u>	<u>03</u>
<u>5.0 Testing Ethernet Networks</u>	<u>04</u>
<u>6.0 Installation and Commissioning</u>	<u>04</u>
<u>7.0 Troubleshooting and Maintenance</u>	<u>04</u>

Understanding the testing requirements and selecting the appropriate solutions to automate and simplify tests as outlined in RFC 2544 is very important. This paper provides valuable information to quickly and accurately turn-up, troubleshoot and maintain Gigabit Ethernet networks.

ABSTRACT

With the advent of gigabit data rates in the already popular Ethernet protocol, this technology is sure to experience phenomenal growth. This increase in data rate enables new, exciting applications for this mature, well-established technology. These applications now demand proper testing and characterization to ensure that the vast amounts of data transported reach their destination successfully and reliably.

Testing in an Ethernet network can be divided into two main categories, those tests which are performed during installation and commissioning and those performed during maintenance and troubleshooting. The major differentiator between these applications is whether the test is active or passive.

For installation and commissioning situations, active tests are employed to test such things as continuity and configuration, throughput, frame loss, back-to-back frames (commonly called burstability) and latency. For these tests, traffic is generated, up to full line rate, to characterize the capability of the network. Once the network's capability is well documented, Service Level Agreements (SLAs) may be established with confidence.

For maintenance and troubleshooting applications, it is generally undesirable to insert additional test traffic into the network. If a customer is already complaining of a slow network, blasting full line rate traffic into the network will only exacerbate the problem. Instead, in these situations it is best to monitor the network and observe the nature of the traffic on the network to gain important insight into the cause of the original problem.

1.0 What is Ethernet?

Simply stated, Ethernet is a networking technology employed to connect computers and computer devices (printers, servers, PCs, etc.). Ethernet provides Layer 2 functionality in the seven-layer OSI model and as such, its main purpose is to provide point-to-point connectivity between these devices. Ethernet is, in no uncertain terms, the dominant LAN technology occupying over 90% of today's enterprise LANs. While Ethernet provides much of the same functionality as other Layer 2 technologies such as SONET/SDH and ATM, it does so in a frame based manner. This means that computers and computer devices communicate via an exchange of information encoded in Ethernet frames.

2.0 Ethernet Frame Structure

The basis of digital communication is, of course, the transmission and reception of ones and zeros. For the proper interpretation of these ones and zeros however, these signals must be properly structured, that is, the digital bits (ones and zeros) comprising an Ethernet frame must adhere to certain rules. The IEEE 802.3 Ethernet standard defines the rules for how an Ethernet frame is structured.

The 802.3 standard states that basic Ethernet frames may be 64 to 1518 bytes in length, not including the Preamble and Start of Frame delimiter (these are defined below). The fields included in a basic Ethernet frame are as follows;

Destination Address:

A 6-byte value specifying the target station for the frame.

Source Address:

A 6 byte value specifying the sending station of the frame.

Length/Type:

A 2-byte value indicating either the Length of the data field (payload) or the Type of protocol that generated the frame.

Data Field:

46 - 1500 bytes of payload. This field contains the user data; the actual information being transmitted. If the data field does not contain the minimum required 46 bytes, a pad is added to ensure the frame meets the basic length requirement.

Frame Check Sequence (FCS):

A 4-byte value used for error detection. Layer 2 devices throughout the network (computers, switches, bridges, etc) recalculate the FCS when they receive the frame. If the FCS calculated by the receiving (or intermediate) station, does not match the FCS calculated and appended at the previous station, the frame is considered to be errored and the entire frame is discarded.

Although, not generally considered to be part of the Ethernet frame structure, the following fields also deserve mention. These include the preamble and the start of frame delimiter. These fields provide the following;

Preamble:

A 7-byte (56 bits) field of alternating ones and zeros (10101010...). The preamble provides a clock recovery signal for synchronization purposes.

Start of Frame:

A 1-byte field of alternating ones and zeros ending with 2 ones. (10101011). When the receiving station receives two consecutive ones, it knows that the Ethernet frame is about to begin. (It then processes the next 6-bytes as the Destination Address, etc.)

Ethernet Frame Structure							
Preamble	Start Frame Delimiter	Destination MAC Address	Source MAC Address	Length/Type	MAC Client Data	Pad	Frame Check Sequence
7 bytes	1 byte	6 bytes	6 bytes	2 bytes	46 to 1500 bytes		4 bytes

Testing in an Ethernet network can be divided into two main categories, those tests which are performed during installation and commissioning; those performed during maintenance and troubleshooting. The major differentiator between these applications is whether the test is active or passive.

3.0 Emerging Applications for Gigabit Ethernet

As stated earlier, 10/100 Mbps Ethernet networks power over 90% of the world's Enterprise LANs today. 10/100 Mbps Ethernet is a mature technology that has been commercially available for decades, however, it has to date been relegated to LAN data transport. This is largely in part to its historical inability to provide the large bandwidth that MAN, WAN, SAN (Storage Area Networks) and CANs (Campus Area Networks) require. With the advent of gigabit data rates, the restriction of Ethernet to the LAN is no longer valid. Gigabit Ethernet applications include interconnecting LANs, providing high speed Internet connections, MAN backbone and access technology, SAN server interconnect and is even providing cost effective, easy to manage server aggregation in Campus Area Networking environments. These new applications require more extensive testing to be performed because these networks are not only carrying more proprietary LAN traffic, but rather large amounts of end customer traffic as well.

4.0 BERT and Ethernet

As detailed earlier, if the FCS sequence calculated by the device receiving the Ethernet frame does not match the FCS calculated and inserted by the sending station, the entire frame is discarded. Herein lies an important differentiator between a frame-based technology such as Ethernet and a bit-based technology such as SONET/SDH. As a result of this difference, tests such as BER (Bit Error Rate) hold less relevance in Ethernet networks. If even one bit error is detected in an Ethernet frame, the entire frame is discarded. Ethernet frames can be between 64 and 1518 bytes in length, meaning that if the errored frame was a 1518 byte frame and the entire frame were discarded, this would be reported as 12,144 bit errors! (1518 bytes * 8 bits) This would obviously provide incorrect test results and entirely misleading BER values. As a result, if BER tests are to be performed in Ethernet networks they can only be provided on a link-by-link basis (a fiber or cable span with no switches or network elements present). This is generally undesirable as it does not provide any indication of how the overall system will perform once the network elements (switches, bridges, etc.) are installed. Although, a BER test may not apply to a frame based technology such as Ethernet, there are many valuable tests that can be performed.

5.0 Testing Ethernet Networks

Testing in an Ethernet network can be divided into two main categories, those tests which are performed during installation and commissioning; those performed during maintenance and troubleshooting. The major differentiator between these applications is whether the test is active or passive.

For Installation and commissioning situations, active tests are employed to test such things as continuity, configuration, throughput, frame loss, back-to-back frames (commonly called burstability) and latency. For these tests, traffic is generated, up to full line rate, to characterize the capability of the network. Once the network's capability is well documented, SLAs (Service Level Agreements) may be established with confidence.

For maintenance and troubleshooting applications, it is generally undesirable to insert additional test traffic into the network. If a customer is already complaining of a slow network, blasting full line rate traffic into the network will only exacerbate the problem. Instead, in these situations it is best to monitor the network and observe the nature of the traffic on the network to gain important insight into the cause of the original problem.

6.0 Installation and Commissioning

When installing a Gigabit Ethernet network, the first tests to be performed often check basic continuity and configuration. These tests ensure that Ethernet traffic injected into the network, possibly even at full line rate, actually makes it to the correct destination. The type of equipment required for these basic tests provides full line rate traffic generation, the ability to monitor transmitted traffic, receive traffic side-by-side and Ping functionality. With these tests, continuity can be verified and proper configuration of switches and routers can be tested.

Once continuity and proper configuration have been established, the network can be characterized if SLAs are to be committed and upheld. To characterize a network in the field requires the ability to quickly and easily measure maximum throughput, frame loss, back-to-back ability (commonly called burstability) and end-to-end latency. Once these parameters have been well documented, SLA parameters may be established with confidence. The procedures for performing these important tests are detailed in a document titled RFC 2544. The RFC benchmarking tests were originally developed for characterizing devices in a laboratory setting, however, they have been adapted and automated for field measurement applications. For instance, the NetTest CMA 5000 Gigabit Ethernet test module has reached the pinnacle of simplicity with a complete, single-button RFC 2544 characterization test.

7.0 Troubleshooting and Maintenance

Once the network has been properly characterized and entered into service, the same test kit used for installation and commissioning should also be able to provide troubleshooting and maintenance functionality. For these applications it is generally not desirable to generate and inject additional traffic into the network. Instead, a passive approach is often taken, first collecting statistics on the overall, global, network performance and then taking a more detailed analysis to pinpoint the offending station or device.

The most commonly heard complaint by bandwidth suppliers and network maintainers is that the network is slow. In these instances, having the proper test equipment can mean the difference in locating the symptom or the cause. To diagnose such a nebulous problem as, “the network is slow”, requires knowledge of the traffic flow and patterns being carried by the network. This would include information such as;

Throughput:

A measure of the current data rate in Mbps being achieved by the network. Additional information on the maximum, minimum and average Throughput is also helpful.

Utilization:

The percentage of the total available bandwidth currently being consumed. If a network link is over utilized (saturated) it will operate inefficiently as transmission frames on the network must wait for resources to become available. For example, if switch buffers are full and yet the switch continues to receive frames, these additional frames will be dropped resulting in a need for frame retransmission. Conversely, if flow control is enabled in the network, the switch will request that the transmitter stop transmitting until switch buffer resources become available. Either of these scenarios degrades network performance.

Number of errored frames:

Although Ethernet itself is a best effort service and does not guarantee transmission success, higher layer protocols such as TCP, will guarantee success by requesting retransmission of errored (dropped) Ethernet frames. This retransmission leads to inefficiency and degraded network performance by consuming additional bandwidth.

The most commonly heard complaint by bandwidth suppliers and network maintainers is that the network is slow. In these instances, having the proper test equipment can mean the difference in locating the symptom or the cause.

Frame size distribution:

If the traffic on the network consists of a large percentage of small frames (64 bytes), a larger percentage of the overall traffic is overhead (addressing information, length/type field, etc.) In fact, if only 64 byte frames were utilized, approximately 30% of the available bandwidth would be consumed by overhead. Comparing this to less than 2% of overhead for 1518 byte frames, clearly illustrates that using larger frames provides more useable bandwidth.

To further assist in pinpointing the source of the problem some test instruments such as the CMA 5000 offer the ability to view statistics on a much more detailed, per channel basis. These channel statistics provide information such as; frames and errors generated per Ethernet or IP address and frames and errors generated per VLAN or MPLS label, among other important values. Having these per channel statistics available, allows technicians to quickly identify which station or VLAN is generating the most traffic or errors. This is a key element to quickly diagnosing and rectifying an over-utilized (a.k.a. slow) network.

Another important application for Gigabit Ethernet test equipment is the ability to verify rate limiting or 'bandwidth throttling'. Gigabit Ethernet networks supply enormous amounts of bandwidth, so much so that some end users may not be able to justify leasing an entire gigabit of bandwidth. In such cases, it is possible to limit the forwarding rate of switches in the network such that the end user may only access 200 or 300 Mbps of bandwidth, or whatever amount they require. Limiting this forwarding rate is referred to as either 'rate limiting' or 'bandwidth throttling'. In this scenario, the service provider may establish a policy within the switch to provide the rate limiting. Once established however, it is almost inevitable that at some point in the network's life, the end user will complain that the network is slow. When this occurs, it is often the burden of the service provider to prove that the agreed upon throughput rate is actually being delivered. Without adequate test equipment, often the only recourse left available to the supplier is to disable the policy in the switch, giving the end user access to the full bandwidth of the gigabit link. While this will obviously rectify the problem, it does present a rather unattractive business case, to say the least. Many times the solution to what appears to be a slow network will have nothing to do with the actual network itself. As previously discussed, frame size distribution, errored frame retransmission and utilization all impact the users experience and their perception of the network's performance.

With the advent of gigabit data rates in the already popular Ethernet protocol, this technology is sure to experience phenomenal growth. This increase in data rate enables new, exciting applications for this mature, well-established technology. These applications now demand proper testing and characterization to ensure that the vast amounts of data transported reach their destination successfully and reliably. Test instruments, such as the CMA 5000 Gigabit Ethernet test module automate and simplify these tests such as RFC 2544 and PING. In addition, the CMA 5000 provides full global statistics and channel statistics to facilitate fast, efficient troubleshooting and maintenance.

To further assist in pinpointing the source of the problem some test instruments such as the CMA 5000 offer the ability to view statistics on a much more detailed, per channel basis. This is a key element to quickly diagnosing and rectifying an over-utilized (a.k.a. slow) network.



NetTest North America Inc.
Center Green, Building 4
6 Rhoads Drive
Utica, NY 13502 USA
Toll Free: 1 800 443 6154
Tel: +1 315 266 5000
Fax: +1 315 798 4038
E-mail: info@nettest.com
Web: www.nettest.com

NetTest Sales Offices

Brazil	+55 11 5505 6688	Italy	+39 02 95 12 621
China	+86 10 6467 9888	Singapore	+65 6220 9575
Denmark	+45 72 11 22 00	Spain	+34 91 372 92 27
France	+33 1 49 80 47 48	USA	+1 315 266 5000
Germany	+49 89 99 89 01-0		

NetTest, the pioneer in multi-layer network testing, is a global provider of test and measurement systems, instruments and components for all types of networks and all stages of network development and operation. Our solutions offer leaders in optical, wireless and fixed networking vital insights into network performance, enabling informed business decisions that drive profitability.