

AN-880 Network Benchmarking



Literature Number: SNLA032

Network Benchmarking

National Semiconductor
Application Note 880
John von Voros
Bonnie Wilson
February 1993



INTRODUCTION

HOW DOES PERFORM3 WORK?

WHAT DELAYS AFFECT THROUGHPUT?

DETERMINING THE PROPER ENVIRONMENT

DETERMINING CLIENT SUITABILITY

HOW DO I CHOOSE A CLIENT SOLUTION?

SERVER SUITABILITY

HOW DO I CHOOSE A SERVER SOLUTION?

CONCLUSION

INTRODUCTION

As the competition for Ethernet® sockets and board level sales has increased, performance has become a means of differentiating a product from the "pack." Most marketers will quote figures given by a benchmark program known as PERFORM3 which is a throughput test provided by Novell® in their driver development kit. The intent of this program is to evaluate the stability of a driver and its associated hardware, not necessarily the throughput. Unfortunately, for lack of any other metric, this has become the de facto standard for the evaluation of performance.

The results provided by PERFORM3 are dependent on the adapter cards, their associated software drivers, the performance of the PC's used (i.e., 80486 machines vs 80286 machines), the number of workstations, and even such things as the length of the interconnecting cable. Therefore, it is relatively easy to skew this data towards whatever result is desired. For example, if you wanted to prove that your client solution was better than another, you might choose to

use a fast PC for the client and a slow PC for the server. Data from this test must be used for relative comparison between cards in the same "fair" environment in order to be useful.

HOW DOES PERFORM3 WORK?

When executing PERFORM3, the user must specify the range of file sizes to be transferred, the step size (i.e., use 1K to 10K file sizes in 1K increments), and the amount of time that each file size will be tested (i.e., 30 seconds). PERFORM3 creates a file of specified length and places that information in the cache memory of the server. This is done to eliminate the delay caused by the hard disk drive. (Hard drives are almost always the rate limiting factor in a server-to-workstation transfer.) During the test interval, each workstation simultaneously requests the cached file. A file transfer is accomplished by requesting multiple reads of a given packet length until the entire file is transferred (shown below). The number of bytes of overhead listed was taken from a protocol analyzer evaluation of a PERFORM3 file transfer.

1. Workstation submits a "read file data request" for a file (57 Bytes).
2. Server sends "read file data reply" which includes requested data + 54 Bytes of overhead.
3. Workstation submits another "read file data request" for the next packet.
4. Server sends "read file data reply" . . .

The data given by this test will give a maximum and average throughput for the specified parameters. Reductions in total bandwidth are caused by software overhead, collisions, the preamble field, and the interframe gap. The maximum attainable throughput for a 1024 byte file can be calculated as follows:

Quantity	Description	Time for Transfer (μ s)
1	Read File Data Request (57 Bytes)	45.6 μ s
1	Read File Data Reply (1K Data + 54 Bytes of Novell/Ethernet Overhead)	862.4 μ s
2	Preamble Fields (8 Bytes)	2 x 6.4 = 12.8 μ s
2	Interframe Gaps	2 x 9.6 = 19.2 μ s
		<hr/>
		940 μ s = <u>1090 Kbytes/sec</u>

Note: Time for Transfer = Number of Bytes / 1.25 Mbytes/sec.

AT/LANTIC™ is a trademark of National Semiconductor Corporation.
Ethernet® is a registered trademark of Xerox Corporation.
Novell® is a registered trademark of Novell, Inc.
3Com® is a registered trademark of 3Com Corporation.

WHAT DELAYS AFFECT THROUGHPUT?

A throughput of 1090 Kbytes/sec represents an ideal network in which there are no collisions or software delays. A more accurate model would take into account all of the relevant delays as shown in *Figure 1*. The delay (D_0) caused by the redirector (TSR program) running on the client is required to intercept BIOS calls that are made for network services. D_1 and D_5 represent the software overhead for the client drivers which provide the software interfaces to the Ethernet cards. In general, the size of the driver in memory is proportional to the amount of delay introduced. Delays D_2 and D_4 are directly related to the efficiency and throughput of the network hardware. D_3 is the cable delay which is a function of the cable length and type. The network operating system introduces a delay in order to accomplish all of its tasks. Finally, D_7 is the delay associated with the server disk drive which tends to dwarf the sum of the other delays. In the event of a collision, the two sending workstations that caused it will wait a random amount of time (determined by the random back-off algorithm specification of I.E.E.E. 802.3) before retransmit. This will introduce a delay due to the time lost while transmitting the collided packet plus the wait for retransmit. Out of these delays, the manufacturer of an Ethernet solution can only control those related to the hardware interface, which is a function of the Ethernet controller chosen, and its software drivers.

DETERMINING THE PROPER ENVIRONMENT

The proper testing environment should use the same types of machines that will be used in the typical end network. In general, most vendors will not use more than 5-10 PCs in their test due to resource/time constraints. This number of workstations does, however, provide enough network traffic to represent a loaded network. For the purpose of this paper, several series of benchmarks will be presented with the environment listed in *Figure 2*. The PERFORM3 parameters specified are:

- Test Time:** 30 Seconds
- File Sizes:** 1K to 10K Bytes
- Increment:** 1K Bytes

Server	
Compaq, Model CP3301, 486DX-50, EISA-Bus, 8MB DRAM	
SONIC-EISA PLX 32-Bit Busmaster Ethernet Card	
Novell Speed Rating: 1327	
Workstations	
1.	AST Premium, Model 5, 486DX-33, EISA-Bus
2.	PC Brand, Model A84310, 486DLC-33 (Cyrix)
3.	PC Brand, Model A84210, 486DX2/50, ISA-Bus
4.	Clone, Western Digital Motherboard Model WDAP4200 (Piranha 4200). ISA-Bus
5.	Dell Model 433DE, 486DX-33, EISA-Bus

FIGURE 2. Test Environment

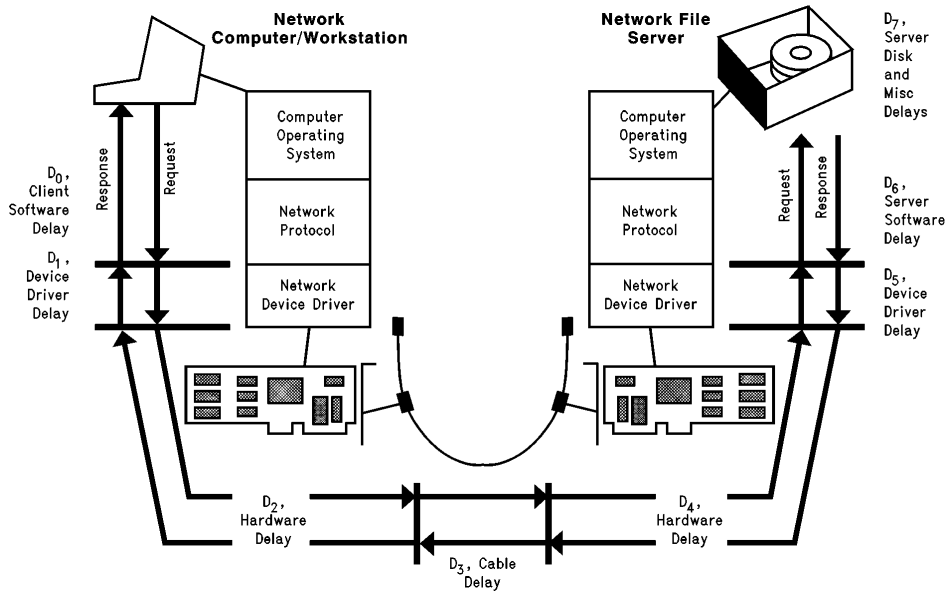


FIGURE 1. Network Delays

TL/F/11791-1

DETERMINING CLIENT SUITABILITY

One of the favorite tests proposed by Ethernet vendors utilizes a high speed server with only one client. It is important for the validity of this test that the server is not the bottleneck in the system. For this benchmark, each client card was tested in the AST machine. It should be noted that EISA machines tend to penalize I/O mapped designs because I/O cycles are slower than in ISA machines.

These results indicate that the 3Com® adapter was the fastest one available, which might be true if you could find a network with only one client. A more realistic test would include some kind of network loading. In order to determine how these boards would perform in a more typical network, five client machines were loaded with five copies of each vendor's card. As shown in *Figure 4*, these adapters all perform within 2.5% of each other when evaluated in a reasonable environment (the raw data is provided at the end of this note). This delta is well within the margin of error, so throughput differences among vendors are slight at best.

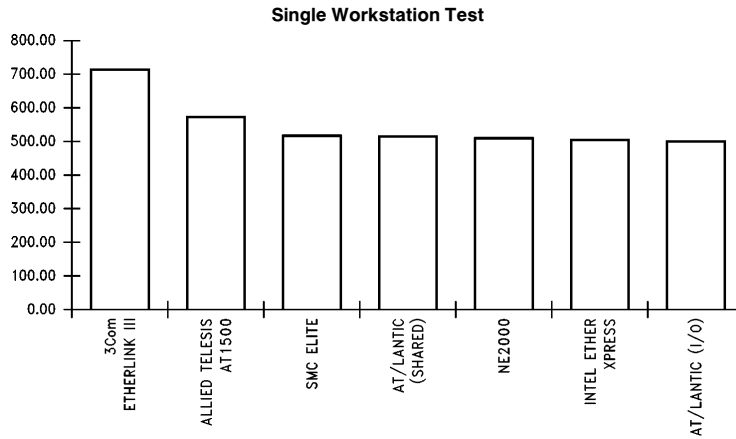


FIGURE 3. Single Workstation, Single Server Throughput Results

TL/F/11791-2

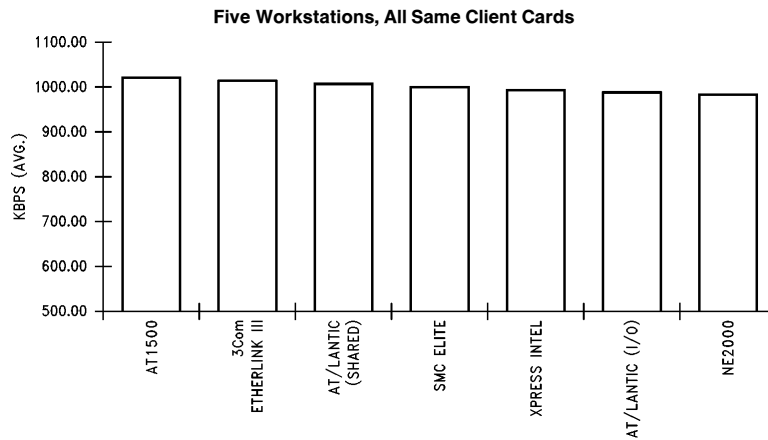


FIGURE 4. Five Workstations, Single Server Throughput Results

TL/F/11791-3

HOW DO I CHOOSE A CLIENT SOLUTION?

Given that all cards perform pretty much the same in a workstation on a loaded network, how does one choose between one supplier or another? This decision is usually based on available software and compatibility with existing hardware. As far as software drivers go, you would like to choose the architecture that shows up as an install option on the menu of the most software packages. That way, there is an extensive installed base of software and driver support provided by the software supplier. Hardware compatibility means that the adapter card or motherboard solution will migrate between different platforms without extra effort. Bus-mastering, while ideal for servers (see next section), tends to have compatibility problems with ISA based machines. The reason for this is that ISA machines were not designed to accommodate bus-mastering devices. Each core logic chip set has different timing parameters which affect the operation of the bus. The I/O and shared memory adapters tend to have the least problems because this is the standard method of interfacing to slave cards on the ISA bus.

SERVER SUITABILITY

Server applications require a different assessment of performance than clients because the server must react to nearly every packet that appears on the network. Clients are only responsible for their small percentage of the overall network load. If a client has low throughput, only the client is affected. A slow server, on the other hand, will lead to a slow network. Another metric to be considered is CPU utilization since a server with no leftover CPU bandwidth may drop packets or be unable to run multiple modules of software.

In order to determine the suitability of different cards for a server, NE2000 cards were placed in the five client machines to represent a constant load. The CPU figures quoted in *Figure 5* represent the CPU utilization given by Novell's MONITOR program running on the server. Although this is a common practice, that figure represents more than just the bandwidth requirements for the transfer of packets. 32-Bit server adapters tend to have less CPU utilization than 16-bit cards due to double work transfers.

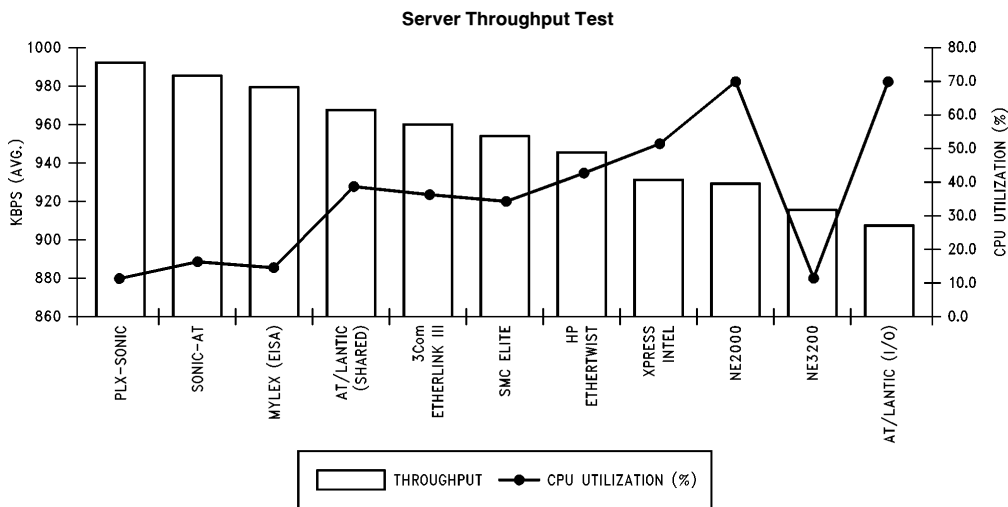


FIGURE 5. Server Throughput and CPU Utilization

TL/F/11791-4

The optimum solution for a server card depends on the type of network that will be used. If a network will only consist of five clients and needs little room to expand, the NE2000 card would be the choice because of its low cost and hardware stability. On the other extreme, if a user needs to put more than a few Ethernet cards in the server to create multiple segments, the NE3200 would be the choice because it has the lowest overall CPU utilization. This card would excel in heavily loaded servers because of an embedded micro-controller that is running the protocol and thus unburdening the host CPU. A useful measure for the average server application would combine performance and CPU utilization into one figure of merit. For the purpose of this evaluation, *Figure 6* illustrates the various cards as ranked by **THROUGHPUT/CPU UTILIZATION** for the data shown in *Figure 5*.

CONCLUSION

In conclusion, Ethernet cards should be evaluated in a realistic environment as opposed to special cases that may highlight certain aspects of the cards. As Ethernet has become a mature technology, the throughput of most cards has approached its bandwidth limitations. Ideally, the end user should run these benchmarks on the target environment to obtain unbiased results. More realistic tests would include scripting of actual file transfers and possibly what effects the card may have on other aspects of the machine such as video performance. It should be noted that a faster hard drive on the server provides the most tangible performance increase to the end user.

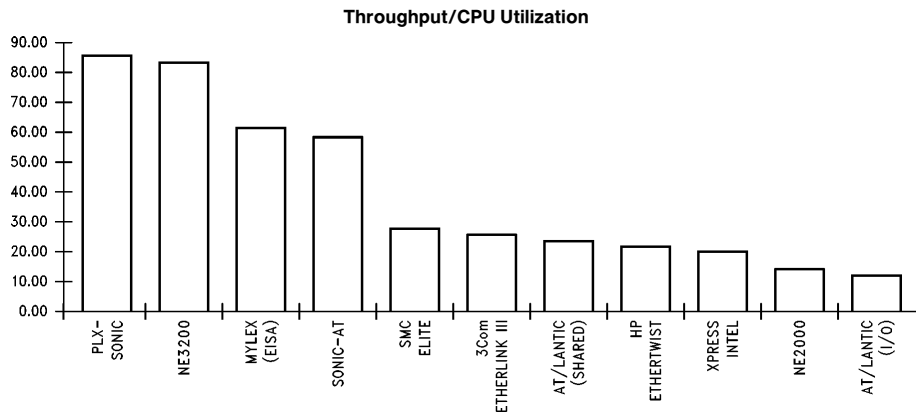


FIGURE 6. Server Throughput/CPU Utilization

TL/F/11791-5

DATA

Figure 3

Card	Avg. Throughput kbytes/S
3Com ETHERLINK III	731.57
ALLIED TELESIS AT1500	578.72
SMC ELITE	519.64
AT/LANTIC (SHARED)	514.77
NE2000	501.46
INTEL ETHER XPRESS	492.00
AT/LANTIC (I/O)	489.39

Figure 4

Card	Avg. Throughput kbytes/S
AT1500	1021.69
3Com ETHERLINK III	1015.16
AT/LANTIC (SHARED)	1005.78
SMC ELITE	1003.27
XPRESS INTEL	999.31
AT/LANTIC (I/O)	998.25
NE2000	996.66

Figures 5 and 6

Card	Avg. Throughput kbytes/S	CPU Utilization (%)	Throughput/ CPU Utilization
PLX-SONIC	992.76	11.5	86.33
NE3200	916.75	11.0	83.34
MYLEX (EISA)	981.49	16.0	61.34
SONIC-AT	986.15	17.0	58.01
SMC ELITE	953.49	35.5	26.86
3Com ETHERLINK III	960.81	37.0	25.97
AT/LANTIC (SHARED)	966.12	38.6	25.03
HP ETHERTWIST	946.59	42.5	22.27
XPRESS INTEL	932.33	50.0	18.65
NE2000	930.91	70.0	13.30
AT/LANTIC (I/O)	907.68	71.6	12.68

LIFE SUPPORT POLICY

NATIONAL'S PRODUCTS ARE NOT AUTHORIZED FOR USE AS CRITICAL COMPONENTS IN LIFE SUPPORT DEVICES OR SYSTEMS WITHOUT THE EXPRESS WRITTEN APPROVAL OF THE PRESIDENT OF NATIONAL SEMICONDUCTOR CORPORATION. As used herein:

1. Life support devices or systems are devices or systems which, (a) are intended for surgical implant into the body, or (b) support or sustain life, and whose failure to perform, when properly used in accordance with instructions for use provided in the labeling, can be reasonably expected to result in a significant injury to the user.
2. A critical component is any component of a life support device or system whose failure to perform can be reasonably expected to cause the failure of the life support device or system, or to affect its safety or effectiveness.



National Semiconductor Corporation
2900 Semiconductor Drive
P.O. Box 58090
Santa Clara, CA 95052-8090
Tel: (800) 272-9959
TWX: (910) 339-9240

National Semiconductor GmbH
Livny-Gargan-Str. 10
D-82256 Fürstenfeldbruck
Germany
Tel: (81-41) 35-0
Telex: 527849
Fax: (81-41) 35-1

National Semiconductor Japan Ltd.
Sumitomo Chemical
Engineering Center
Bldg. 7F
1-7-1, Nakase, Mihama-Ku
Chiba-City,
Ciba Prefecture 261
Tel: (043) 299-2300
Fax: (043) 299-2500

National Semiconductor Hong Kong Ltd.
13th Floor, Straight Block,
Ocean Centre, 5 Canton Rd.
Tsimshatsui, Kowloon
Hong Kong
Tel: (852) 2737-1600
Fax: (852) 2736-9960

National Semicondutores Do Brazil Ltda.
Rue Deputado Lacorda Franco
120-3A
Sao Paulo-SP
Brazil 05418-000
Tel: (55-11) 212-5066
Telex: 391-1131931 NSBR BR
Fax: (55-11) 212-1181

National Semiconductor (Australia) Pty. Ltd.
Building 16
Business Park Drive
Monash Business Park
Nottingham, Melbourne
Victoria 3168 Australia
Tel: (3) 558-9999
Fax: (3) 558-9998

National does not assume any responsibility for use of any circuitry described, no circuit patent licenses are implied and National reserves the right at any time without notice to change said circuitry and specifications.

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

TI products are not authorized for use in safety-critical applications (such as life support) where a failure of the TI product would reasonably be expected to cause severe personal injury or death, unless officers of the parties have executed an agreement specifically governing such use. Buyers represent that they have all necessary expertise in the safety and regulatory ramifications of their applications, and acknowledge and agree that they are solely responsible for all legal, regulatory and safety-related requirements concerning their products and any use of TI products in such safety-critical applications, notwithstanding any applications-related information or support that may be provided by TI. Further, Buyers must fully indemnify TI and its representatives against any damages arising out of the use of TI products in such safety-critical applications.

TI products are neither designed nor intended for use in military/aerospace applications or environments unless the TI products are specifically designated by TI as military-grade or "enhanced plastic." Only products designated by TI as military-grade meet military specifications. Buyers acknowledge and agree that any such use of TI products which TI has not designated as military-grade is solely at the Buyer's risk, and that they are solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI products are neither designed nor intended for use in automotive applications or environments unless the specific TI products are designated by TI as compliant with ISO/TS 16949 requirements. Buyers acknowledge and agree that, if they use any non-designated products in automotive applications, TI will not be responsible for any failure to meet such requirements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

Products

Audio	www.ti.com/audio
Amplifiers	amplifier.ti.com
Data Converters	dataconverter.ti.com
DLP® Products	www.dlp.com
DSP	dsp.ti.com
Clocks and Timers	www.ti.com/clocks
Interface	interface.ti.com
Logic	logic.ti.com
Power Mgmt	power.ti.com
Microcontrollers	microcontroller.ti.com
RFID	www.ti-rfid.com
OMAP Mobile Processors	www.ti.com/omap
Wireless Connectivity	www.ti.com/wirelessconnectivity

Applications

Communications and Telecom	www.ti.com/communications
Computers and Peripherals	www.ti.com/computers
Consumer Electronics	www.ti.com/consumer-apps
Energy and Lighting	www.ti.com/energy
Industrial	www.ti.com/industrial
Medical	www.ti.com/medical
Security	www.ti.com/security
Space, Avionics and Defense	www.ti.com/space-avionics-defense
Transportation and Automotive	www.ti.com/automotive
Video and Imaging	www.ti.com/video

TI E2E Community Home Page

e2e.ti.com

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2011, Texas Instruments Incorporated