



Evaluating Network Adapters: Methodology and Analysis

January, 1996

Table of Contents

Where to Start	2
New Performance Technologies	2
Evaluation Considerations	3
Performance Testing How-to's	7
Appendix A. Recording Your Evaluation Results	12
Appendix B. Recording Your Features Comparison ...	13

Cogent develops the *eMASTER+* series of Ethernet and Fast Ethernet adapters recognizing that a network is more than just a server, a workstation, and a benchmark. The company's unique systems approach to Ethernet and Fast Ethernet networking targets complex, heavily loaded corporate networks running resource-intensive applications.

Many adapters are designed to perform well under the isolated conditions of typical benchmark tests, but their performance may not hold up in high traffic environments indicative of your network. Likewise, while you may find many adapters in your price range, not all will offer the features you need. This document will help you identify the adapters that will best meet all of your selection criteria.

Where to Start

Conducting network adapter evaluations can be a daunting task, so we will take a step-by-step approach. First, we will discuss two new technologies that have made a huge impact on network performance expectations. Then, we will discuss some of the technical factors you should consider in your evaluation. Finally, we will give you some practical how-to's on performance testing, and will discuss a pair of widely used performance testing tools.

New Performance Technologies

Two major technologies have surfaced during the 1990s that significantly impact price/performance expectations. These two technologies are a new network specification known as Fast Ethernet, and a new PC bus design, known as Peripheral Component Interconnect, or "PCI."

Fast Ethernet

Ratified by the Institute of Electrical and Electronic Engineers (IEEE) on June 14, 1995 under IEEE 802.3u, Fast Ethernet is a Carrier Sense Multiple Access Collision Detect (CSMA/CD) network designed to move data at the rate of 100 Megabits, or 12.5 Megabytes (100 Megabits divided by 8 bits per byte), per second. This data throughput rate represents a tenfold increase over the speed of IEEE 802.3 standard Ethernet, which runs at 10 Megabits, or 1.25 Megabytes, per second. And, as you might surmise from the specification numbers of Ethernet and Fast Ethernet (802.3 and 802.3u), Fast Ethernet is a derivation of standard Ethernet.

The IEEE further divided Fast Ethernet into three sub-categories, based on the type of cabling used:

IEEE 802.3u 100BASE-T4. This version of Fast Ethernet runs on 4 pair of CAT 3 UTP cabling, and can also be run on CAT 4, or CAT 5 UTP, as well as Type 1 STP.

IEEE 802.3u 100BASE-TX. This version of Fast Ethernet runs on 2 pair of CAT 5 UTP cabling, and can also be run on Type 1 STP.

IEEE 802.3u 100BASE-FX. This version of Fast Ethernet runs on 62.5/125 micron multimode fiber optic cabling (the same type of fiber used for FDDI and 10-BASE-FL).

The significance of Fast Ethernet to network performance expectations is obvious—network managers must now analyze the costs and benefits of their own installations against the costs and benefits of running Fast Ethernet.

The PCI Bus

The next performance technology to be considered, the Peripheral Component Interconnect (PCI) bus, has become the dominant PC architecture over the last two years. Prior to the advent of the PCI bus, the fastest personal computer bus available was the EISA bus, which provides a theoretical maximum of 33 Megabytes per second, and a practical maximum of 8 Megabytes per second. In comparison, the PCI bus provides 132 Megabytes per second of bandwidth. Especially in the context of the bandwidth necessary to support Fast Ethernet, PCI plays a critical role in overall system performance.

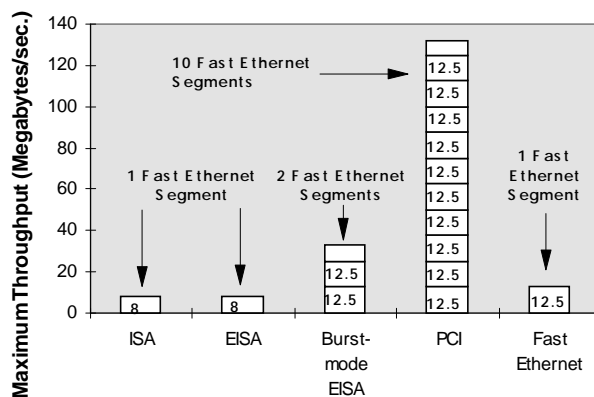


Figure 1. Technology Throughput Comparison

Bus speeds are often translated into their Ethernet or Fast Ethernet equivalent by evaluators (see Figure 1). Keeping that in mind, the average EISA bus (8 Megabytes per second) provides enough bandwidth for approximately one Fast Ethernet segment. The PCI bus, in contrast, provides bandwidth sufficient to support ten Fast Ethernet segments. You should consider, too, that your system's

bus must provide bandwidth for more than just Ethernet or Fast Ethernet segments, supporting SCSI and IDE disk devices on servers, and high-performance video devices on workstations. As those devices become more bandwidth intensive, with the development of Fast SCSI and video compression techniques, network designers need to make sure that their system bus keeps up with the technologies being plugged into them.

Evaluation Considerations

When testing networks for performance, you should consider that many factors impact final results. The nature of the traffic being placed on the cable, traffic management algorithms and adapter traffic handling techniques, and the speed with which an adapter can obtain use of the cable all play a part in overall system performance.

The Nature of Network Traffic

Although applications exchange messages, the currency of workstations and servers is packets. In order for an application such as Microsoft Word to send a message across a network, the message must typically be broken into multiple packets of data which can be sent to a workstation or server on the other end of the cable. In both the Ethernet and Fast Ethernet world, packets sent between a workstation and server contain between 64 bytes and 1,514 bytes, or 1.5 Kilobytes (KBytes) of data. Data sent across networks requires packets of varying sizes; data greater than 1.5 Kbytes is broken into 1.5 Kbyte packets while data smaller than 64 bytes is padded to fit into 64 byte packets. For example, an application on a workstation might send a request to a server to open a file. If the message to fulfill this request contains 6,096 bytes of data, it is packetized into four 1.5 Kbyte packets and one 64 byte packet (40 bytes plus 24 bytes of padding) for transmission on the cable to the server. Note that both Ethernet and Fast Ethernet traffic adheres to these same principles.

In the past, network communications were characterized by text-based applications such as word processing and spreadsheet programs which generally send small messages requiring only a few packets. As networks have become more dependent on resource-intensive applications such as Windows, graphics, and multimedia, the characteristics of network communications has changed; the large file sizes of these applications require large messages and more packets. While a typical spreadsheet file might contain 10 Kilobytes (KBytes) of data, the average graphics file might contain several *Megabytes* (MBytes) of data.

Traffic across networks does not consist just of data packets, it also includes control packets which establish

the communication links between workstations and servers. Control packets contain information such as requests and acknowledgments that prepare the way for transmission of the actual data sent across the network. Much of this communication between workstations and servers is added overhead to the actual transmission of data. Network operating systems generally have a high ratio of control packets to data packets. With the new graphics-based applications and large data message sizes, large numbers of control packets result in unacceptably high overhead and low performance. Changes in applications and network operating system designs have changed the nature of traffic common on today's networks. The following chart contrasts the distribution of packets on typical networks in the past and present.




Figure 2. Typical Packet Size Distribution

Traffic Management Algorithms

Many vendors have recognized that sending excessive numbers of control packets is inefficient and negatively impacts networking performance. As a result, they have developed their newest versions of network operating systems to use fewer control packets to accomplish the same tasks. Instead of sending an acknowledgment after each packet is received, many network operating systems now require an acknowledgment only after multiple packets have been received.

As one example of intelligent traffic handling, Novell has implemented 'packet bursting' technology in its DOS VLM client software and its NetWare 3.12 and 4.x servers. This technology allows the use of a single acknowledgment in response to the reception of multiple packets, rather than requiring one acknowledgment per packet received. The latest VLM is downloadable freely from Novell's !GO NOVFILES forum on CompuServe, as VLMUP3.EXE. NetWare 3.12 and 4.x servers implement packet bursting transparently and automatically. As another example, TCP/IP, a popular protocol for Windows NT, utilizes sliding windows (which allows variability in the number of data packets delivered before an acknowledgment is required) to decrease the ratio of control packets, effectively increasing the throughput of the system.

With the release of NetWare 4.x, Novell has also introduced Novell Link Services Protocol (NLSP), replacing the older Routing Information Protocol (RIP)/Service Advertising Protocol (SAP) protocol. NLSP increases effective throughput by reducing the overhead created by servers' broadcasting routing and service information to the network. With NLSP, each server broadcasts routing information and service information only at bootup time, and when any change occurs in routing or services being provided. NLSP is downloadable freely from Novell's !GO NOVLIB CompuServe forum as IPXRTR.EXE, and can be installed on any NetWare 3.12 or 4.x server.

As the ratio of network control data to user data (spreadsheets, word processing documents, etc.) decreases, the throughput rate (representing the rate at which user data is transferred) actually increases. So, regardless of the network operating system you are using, make sure that you are taking advantage of all of the features it has to offer.

Traffic Management Techniques Employed by Adapters

In addition to protocol-based technologies already mentioned, adapters also employ methods intended to increase effective data throughput. Some technologies, such as 3Com's Parallel Tasking and SMC's SimulTasking, take aim at smaller packets and networks with lower traffic loads. Given the changing nature of network traffic, Cogent has chosen to focus specifically on large message handling and busy networks with its eMASTER+ adapters, offering Predictive Pipelining for Novell NetWare environments.

Traffic Handling on Standard Ethernet and Fast Ethernet Adapters

Data transmission across Ethernet, or Fast Ethernet, networks involves more than just throwing data onto a wire and pulling it off again on the other side. Many steps are necessary to prepare the data for transmission and send it onto the cable. Data sent across networks must be processed at the application, network operating system (NOS), network driver, and hardware levels. With each step in the process adding overhead, inefficient designs and bottlenecks can significantly impact network performance. Many standard Ethernet adapters require intensive data processing with numerous memory-to-memory data copies, data queuing, and interrupts—all of which slow the delivery of data in the network.

For example, in standard Ethernet adapters, a simple task, such as loading a word processing file from a network server, requires breaking the message from the application into packets the driver can handle, calling the driver for

each packet, copying each packet from network operating system (NOS) memory to network driver memory, queuing each packet to be sent on the cable, and so on. In fact, it might surprise you that the time data actually spends on the cable is only about 15% of the total data transfer time. With Fast Ethernet, this percentage is even smaller.

Although this method of processing data was adequate in the days of small messages, you can imagine what happens when a standard Ethernet adapter handles a large message typical of a graphics-based application. Every interrupt, copy, and queued packet negatively impacts efficiency and performance. For a relatively small 64 KByte message, multiply this high overhead 48 times (64 KBytes divided by 1.5 KBytes per packet) and you get the idea.

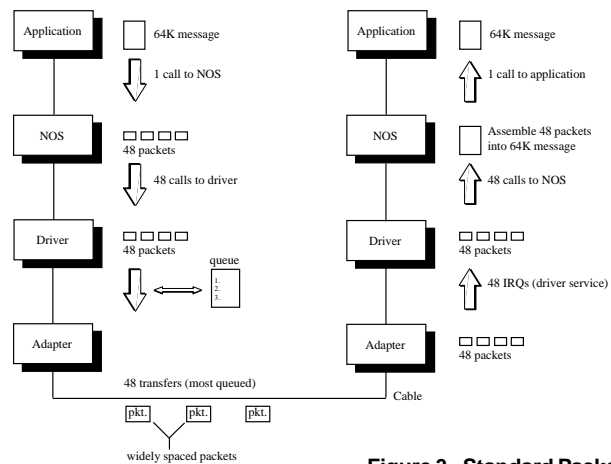


Figure 3. Standard Packet Processing

Traffic Handling on Third-Generation Ethernet and Fast Ethernet Adapters

Some vendors have attempted to attack the problem of overhead with new third generation technologies such as 3Com's Parallel Tasking and SMC's SimulTasking. These technologies focus primarily on overhead on the receive side and do not affect the transmit side. Data processing on the receive side in parallel tasking adapters is slightly more efficient due to the overlapping of driver execution time and packet reception time. However both the sending side and the receiving side still require the same number of operations as standard Ethernet adapters—as shown in Figure 3.

Furthermore, these third generation parallel tasking technologies suffer significant performance losses in heavily loaded networks. For example, 3Com's parallel tasking technology relies on an early interrupt (IRQ) that allows the driver to notify the NOS of an incoming packet before the packet has been completely received. If, however, the packet isn't received, perhaps due to a collision on the cable, the driver and NOS must go through a high overhead recovery process which degrades

performance in loaded networks. Analysis of these new parallel tasking technologies as a whole reveals that the advantages in receive efficiency with background load is offset by the loss in performance.

General Traffic Handling in Cogent eMASTER+ Adapters

For all types of network traffic, Cogent adapters offer technological advantages that benefit both the sending and receiving side of data communications, and maintain performance even during periods of heavy traffic load. For example, Cogent's PCI Fast Ethernet adapters are equipped with 2 Kilobytes of transmit and 4 Kilobytes of receive high speed First In First Out (FIFO) memory. Such large buffers mean that the adapter can buffer traffic during periods of heavy load, maintaining a characteristic even performance profile throughout. Because both incoming and outgoing traffic is buffered, both sides of the transmission process benefit. Cogent's ISA and EISA Fast Ethernet adapters are equipped with 128 Kilobytes of packet buffer space, and so benefit from the same concept.

All of Cogent's adapters implement the design architecture most suited to high performance on their given bus. Cogent's PCI adapters are bus mastering, and so affect data transfers directly from the network cable to system memory, with no time-consuming intermediate data copies between driver and application memory space. Cogent's EISA adapters support 33 Megabytes/sec. burst mode transfers, the fastest possible transfer rate across any EISA bus. Cogent's ISA adapters use an I/O mapped architecture, so that data may be transferred through I/O addressing space, rather than through slower Upper Memory Block (UMB) addressing space.

NetWare IPX Traffic Handling in Cogent eMASTER+ Adapters

For Novell NetWare, the most widely-used network operating system, Cogent designed its innovative Predictive Pipelining technology to dramatically reduce processing overhead by handling whole messages instead of just single packets. Instead of the NOS creating individual packets from a message and calling the driver each time a packet is finished, Predictive Pipelining network drivers hand the whole message to the hardware level, for packetization. Speed and efficiency of data transfers are further enhanced by the predictive capability of the network drivers. Before it receives the whole message, the network driver can predict the total amount of data required for the whole message, not just single packets, and allocate resources in a single step. With fewer steps required to process data on both the transmit

and receive sides (as shown in Figure 4), eMASTER+ significantly reduces unnecessary overhead resulting in substantial performance advantages.

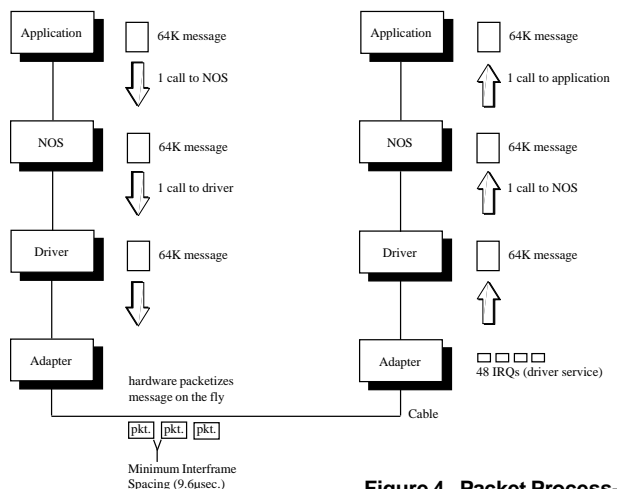


Figure 4. Packet Processing with Predictive Pipelining

Predictive Pipelining also has the advantage of reducing the sensitivity of eMASTER+ adapters to collisions typical of heavy network traffic. Unlike the parallel tasking technologies previously mentioned, Predictive Pipelining does not require lengthy operations between the driver and NOS to handle packets which have collided on the cable. With Predictive Pipelining, only when the whole message has been received does the driver notify the NOS and transfer the entire de-packetized message to the NOS. If a packet experiences a collision on the cable, Predictive Pipelining network drivers simply wait until it is retransmitted rather than aborting the entire message.

The result of Predictive Pipelining is increased data throughput for workstation adapters, and server adapters that can handle many more workstations with lower server CPU utilization.

Interframe Spacing

Network performance is also affected by interframe spacing—the spacing of packets on the cable. The longer it takes for a packet to be sent, the wider the space is between packets on the cable. eMASTER+ adapters actually achieve and maintain Minimum Interframe Spacing of 9.6 µsec. (microsecond), for Ethernet, and 0.96 µsec. for Fast Ethernet. For most adapters with costly overhead and limited buffering capability, spacing between packets on the cable is typically 40 to 140 µsec. for Ethernet, and 0.40 to 14 µsec. for Fast Ethernet.

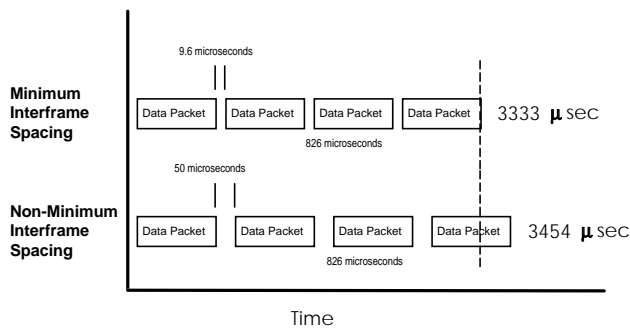


Figure 5. Interframe Spacing

Other Comparison Criteria

In addition to performance, a variety of other features will be important to your evaluation. We've included a chart in Appendix B to to help you weigh these important factors.

Automated configuration.

Does the adapter support PCI BIOS configuration methods (for PCI adapters)? Does the adapter support Plug and Play (for ISA adapters)? Is the adapter configurable by standard EISA configuration tools (EISA adapters)?

Diagnostics.

Is the adapter supplied with a diagnostics tool? Does the diagnostics tool support both local and remote (echo server) tests?

Comprehensive network drivers.

Does the manufacturer provide the drivers you need, for all adapters and all platforms?

Multi-threaded drivers.

Does the adapter come with 32-bit multi-threaded drivers?

Cogent's eMASTER+ hardware designs support multi-threaded software drivers which allow several processes to occur simultaneously in the protected mode of multitasking operating systems such as OS/2, UNIX, and NetWare. Because cable time overlaps processing time, multi-threaded drivers provide significant performance advantages over single-threaded drivers.

Cabling options.

Does the adapter support multiple line speeds (10 and 100 Megabits)? If you only need one speed, can you buy the adapter with just that line speed only, at a cost savings? Does the adapter provide the cable connection you need? If you only need one connector type, can you buy the adapter with just that connector, at a cost savings?

Depending on bus type and line speed, Cogent offers its adapters in both single media and combo (UTP, BNC, AUI) models.

Multipoint adapters.

Does the manufacturer offer multipoint (multisegment) versions of their adapters for servers?

Cogent is the first and only manufacturer of multipoint Quartet adapters for both 10 and 100 Megabits/sec. servers.

Quality of Construction.

Does the adapter come with a lifetime warranty? Does the manufacturer test each adapter immediately before shipment? How long has the manufacturer been in business?

Because Cogent's reputation depends on the quality of its products, all eMASTER+ adapters must meet stringent quality assurance guidelines before they go to market. For example, the eMASTER+ adapters must first pass four independent testing cycles performed during manufacture. In addition, Cogent retests all adapters individually before final shipment to customers. As a result of this quality assurance program, Cogent products have consistently met the high standards of Original Equipment Manufacturers such as Siemens Nixdorf, Intel, Network General, UNISYS, and others who demand extensive testing, proven standards compliance, rigorous quality assurance procedures, and guaranteed lifetime reliability.

Quality Components.

Is the adapter built with market-proven components?

Cogent's adapters are built around the following Ethernet and Fast Ethernet Controller chips.

Intel 82596 and 82595 controller chips. Cogent's eMASTER+ EM525 AT, and EM93x series EISA adapters use Intel's state-of-the-art controller chips, which provide embedded hardware packetization and link list handling to increase the speed of data processing.

DECchip controllers. Cogent's eMASTER+ PCI adapters were the first to use DEC's industry-leading PCI technology chips, both for standard Ethernet and Fast Ethernet. The larger buffer sizes, and Full Duplex capability are two reasons the DEC chips lead the industry in PCI technology.

DEC 21050 controller chips. Used by Quartet adapters, this specialized DEC chip provides a bridge between the four separate PCI Ethernet segments on a Cogent Quartet adapter and the PCI bus on the system motherboard.

SMC FEAST controller chips. Cogent brings Fast Ethernet technology to the ISA and EISA buses with the SMC FEAST chip. These chips provide the features

necessary for 10/100 support, Full Duplex support on both line speeds, Plug and Play support, and I/O mapped architecture to prevent RAM consumption.

Where are the adapters manufactured?

Cogent manufactures all of its products in the U.S.A.

Performance enhancement features.

What performance enhancement features are built into the adapter? Do these features address specifically your network concerns (heavy traffic, large packets, security, cabling distance)?

Performance Testing How-to's

With a knowledge of the context of adapter performance testing, we can now discuss the actual testing.

As much as possible, you should try to test the capabilities of adapters by emulating real network conditions, meaning large numbers of workstations, multiple server segments, simultaneous network activities, and a traffic profile indicative of your own network. You should seek those adapters that provide significantly higher throughput on the cable and lower CPU utilization at the server, even with the cable loaded with background traffic, workstations, and multiple network segments.

General Performance Testing Guidelines

Regardless of the operating system you choose, here are some testing thoughts to keep in mind.

Test within a "clean" environment.

Set up a test network specifically for performance testing so that you can control all of the variables that impact performance. This means that no traffic, other than your testing traffic, should be on the network. Otherwise, variable environmental factors may impact performance numbers and invalidate comparisons between adapters.

Test performance under different network conditions.

Run benchmark tests in your test network both with and without background traffic to measure how well your existing adapters and your new eMASTER+ adapters perform in real world networking conditions. While many adapters perform well without background load, their performance degrades dramatically with other traffic on the network. Cogent eMASTER+ adapters maintain much of the same performance, regardless of network load.

Only make comparisons between data sets run under exactly the same conditions.

Because performance testing is hardware dependent, you should only compare performance data between adapters

generated in precisely the same environment. And, although reproducing testing scenarios published in trade journals can be a useful exercise, don't try to compare their results to your own. Results of performance testing depend on many variables, including the type of server, type of workstation, and amount of memory. Your numbers will almost certainly differ from the published values, even if you duplicate the published test exactly.

Isolate the component.

As much as possible, you want to perform tests that measure variations in the component you are testing. In other words, when testing network adapters, you want to actually measure differences in network adapters, and not differences in hard drive performance.

Test at a relevant file size.

Network performance is very dependant on the size of file being transfered. You should run tests that reflect the median file size transfered on your network.

Keep good records.

The importance of recording actual performance test results, including throughput and CPU utilization, is very apparent. In addition, you will want to record for future reference the date of testing, the type and processor of machines used, driver dates and versions, and operating system and versions.

Performance Testing Tools

The tool you choose for testing will heavily determine the validity of your results. With that fact in mind, here are some general rules to consider about testing tools.

Use a tool that tests the relevant component.

If you are testing network cards, you want to run a test that measures network card performance. For this reason, Norton Speed Index, which measures disk access time, does not do a good job of isolating network adapter performance.

Use a widely distributed and accepted tool.

Your results will be more meaningful if the data generated is readily understood and accepted by others. When talking about adapter performance with customers or resellers, numbers from a standard performance tool will be received more readily than those generated using an exotic methodology. The most widely used performance tools are available free.

The two most widely used network adapter performance testing tools are Perform3 version 1.61, from Novell, and NetBench 3.0 (and the recently released version 4.0), from Ziff-Davis labs. We won't try to document the use of these tools here, since documentation comes with each of the tools. We will give a brief description of each tool,

and where it can be downloaded. Keep in mind that, even if you choose another performance testing tool, the principles of performance testing outlined in this document still apply.

How They Work

Both Perform3 and NetBench run as executables from one or more network clients. The network client sends FILE OPEN and FILE READ requests to the server, and the server responds to these requests by opening the file and replying with the appropriate data. The performance tools ensure that the data which is read from the server exists in the server's cache—a critical issue for performance testing. Consider the following formulas.

test run using uncached files

Total Throughput = Hard Drive Controller Processing Time + Network Adapter Processing Time

test run using cached files (Perform3 and NetBench)

Total Throughput = Network Adapter Processing Time

Since the client's requests are processed from cache, rather than from the server's hard drive, the time required to process the requests can all be attributed to network adapter performance; the tests successfully isolate network adapter subsystem performance from hard drive controller performance. A final feature of both of these tools is that they can run tests using a variety of file sizes. If most of the files read across your network are 64KBytes, then testing file sizes of 2 KBytes is not very useful to you. For this reason, Perform3 can generate requests for files up to 64 KBytes in size; NetBench can generate requests up to 32 KBytes in size. Keep in mind that, as file size increases, adapters are able to take advantage of economies of scale, so throughput increases with file size.




Figure 6. Performance and File Size

Perform3

Perform3 version 1.61 runs as a DOS application on any DOS network client. The executable, PERFORM3.EXE, is stored on a server drive, but is executed from the workstation. Although a utility written by Novell, Perform3 can be used with any network operating system to which a DOS workstation can be connected, and is not

protocol-specific. Perform3 can be run on a single workstation to measure the workstation's adapter throughput; it can also be run simultaneously on a number of workstations to measure a server's network adapter throughput and CPU utilization. Perform3 can be downloaded from the !GO NOVUSER forum in CompuServe, as PRFRM3.ZIP. Its close cousin, Perform2 (not discussed here), can be downloaded as PRFRM2.ZIP.

Perform3, as a command line utility, is the simplest performance tool to use. It allows you to choose several test parameters, including file size, step size, and test duration, at runtime. These parameters greatly impact the precision of your test results. Smaller file step sizes, or increments, show greater variation than larger step sizes, but require significant time to run the tests.

We recommend you test your adapters with 4 KBytes (4096 bytes) file size with 1 KBytes (1024 bytes) step starting at 1024 KBytes for a duration of 12 seconds; type the following at the command line where *filename* is the name of the file that will contain your Perform3 results:

```
perform3 filename 12 1024 4096 1024
```

In tests between one workstation and one server, Perform3 reports the local maximum and the local average as the throughput at the workstation. In tests with multiple workstations, Perform3 reports the local maximum, the local average, the aggregate maximum and the aggregate average. The aggregate results reflect server throughput. Figure 7 shows a sample Perform3 screen.

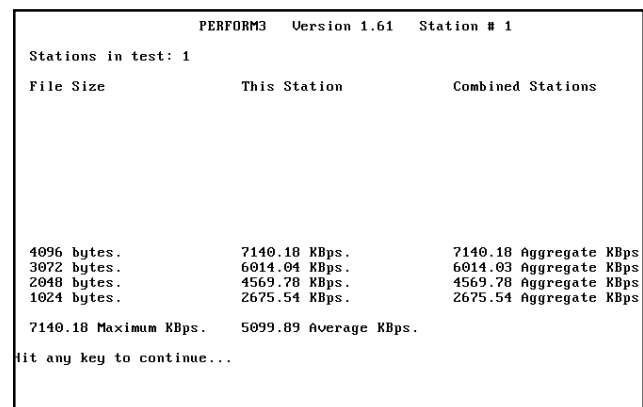


Figure 7. Typical Perform3 Output Screen

NetBench

With the release of Ziff-Davis Lab's NetBench 4.0, NetBench clients are now available for DOS, Windows for Workgroups and Windows NT. Like Perform3, NetBench can be run on a variety of operating systems and protocols, and can be run on single and multiple workstations. Unlike Perform3, NetBench is not limited to DOS clients, but does require a dedicated workstation

to act as a ‘controller’ during tests. NetBench 4.0 does generate nice results files, and, if you have Microsoft Excel installed on the controller workstation, supports macro links to Microsoft Excel, for viewing and formatting results. A close cousin of NetBench is ServerBench, which conducts measurements on other server subsystems, including disk drive and memory.

NetBench 4.0 requires more setup and explanation than Perform3, and substantial documentation is available. NetBench 4.0 is downloadable from Ziff-Davis’s !GO ZNT:ZDBENC on CompuServe as NETBEN.EXE. The documentation is downloadable as NBD0C.EXE.

What to Measure

A few key variables can provide a good forecast of an adapter’s success or failure in your network. PC Week provides a good model for recording and interpreting performance data in their regular NIC Index feature. The PC Week NIC Index feature reviews the performance of network adapters on a regular basis. The NIC Index results published by PC Week on November 13, 1995, for Fast Ethernet adapters shows the Cogent eMASTER+ EM110 10/100 to be the industry leader in two key categories, the Performance/Efficiency Index, and the Price/Performance Index. These indexes, among others, are explained below. Appendix A provides you with sample data in the NIC Index format, as well as a blank chart for recording your own results.

Throughput is the rate at which data is transferred from one point to another, and is usually measured in terms of either Kilobytes, or Kilobits, per second. In the context of network adapters, throughput can be measured at the server or at the workstation.

Workstation throughput measures the data transfer rate of a single workstation communicating to a single server. To find the maximum throughput rate for a workstation, the workstation is attached to a server not currently connected to any other workstations.

Server throughput measures the aggregate data transfer rate of a single server to several clients. Because of the more robust nature of server operating systems, dedicated network servers outperform single workstations by a large amount. So, to find the maximum transfer rate a server’s network card can attain, several workstations must be attached.

CPU utilization, measured in percentage points, measures the amount of time a server’s processor spends servicing the network adapter. CPU utilization is generally only measured on a server, and measuring methods vary. Novell offers a NetWare Loadable Module (NLM), called

STAT.NLM, that measures CPU utilization on NetWare servers. STAT.NLM is downloadable from the !GO NOVLIB forum in CompuServe as STAT.ZIP. (Be aware that running STAT.NLM actually increases CPU utilization by one to two percent.) NetWare’s MONITOR.NLM can also be used to measure CPU utilization, though this requires an individual to constantly watch MONITOR’s reporting screen.

A creature of PC Week’s NIC Index methodology, the *Performance/Efficiency Index* of an adapter is its throughput (measured in Kilobits per second) divided by CPU utilization (in percentage points). This index determines the efficiency with which an adapter provides throughput for the amount of (finite) CPU resources it requires.

Another construct of PC Week, the *Price/Performance Index* factors price into a separate network adapter rating index. While the aforementioned Performance Index can be considered absolutely objective, the Price/Performance Index is more problematic. Price can be interpreted many ways, including manufacturer suggested retail (MSR), reseller, and street. The most objective way of generating a Price/Performance Index is to use a single source for all price quotes. Even so, these prices may not take into account quantity discounts and key reseller competitive price programs, so factor this index into your decision-making process with a grain of salt.

Interframe spacing, the amount of time between back-to-back packets sent from a workstation, is measured in microseconds, or, millionths of a second. The minimum value allowed by the IEEE is 9.6 microseconds for Ethernet and 0.96 microseconds for Fast Ethernet. An adapter with good performance characteristics will maintain interframe spacing at or near those levels. This variable is not easily (or usually) measured. If you have the equipment, though, we encourage you to take a look at this important variable. You will need a hardware-based Protocol analyzer, such as Network General’s Sniffer, that supports the line speed (10 or 100 Mbps) you are using. A software-only analyzer, like Novell’s LANalyzer for Windows, cannot measure interframe spacing.

Test Network Configurations

We are now ready to put together some test network configurations based on all of the information we’ve already discussed.

For Measuring Workstation Throughput

This simple test scenario requires only a single workstation (preferably of the same make/model used on your

network) and a single server. For testing Fast Ethernet, we recommend using 486 or Pentium-based workstations and servers.

By running Perform3 or NetBench between a single workstation and server, you can gauge an adapter's raw throughput. The predictive nature of this result is suspect, however, because the test configuration in no way




Figure 8. Measuring Single Workstation Throughput

resembles most network environments.

For Measuring Workstation Throughput on a Loaded Network

You also want to know how your workstation behaves on a network supporting multiple concurrent data conversations. Because performance tests do not usually provide a mechanism for generating 'background' traffic (traffic unrelated to the actual test packets), you will need to provide background activity yourself. Add two additional peer-to-peer network workstations (i.e., Windows for Workgroups) to the test scenario described above. From these new workstations, generate background traffic typical of a heavily loaded network by running a batch file which copies files from one of the new workstations to a NUL file on the second new workstation. The following is an example of a batch file you can create using a text editor:

```
:DOAGAIN
COPY F:\WINDOWS\*. * NUL
GOTO DOAGAIN
```

You can run this batch file from one of the peer-to-peer workstations (after making the necessary network

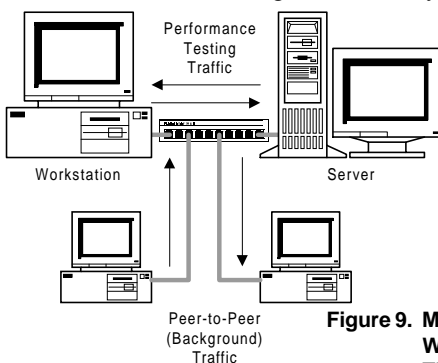


Figure 9. Measuring Single Workstation Throughput with Background Traffic

redirection, in this case to drive *f*) to generate background traffic.

Or, if you have the advantage of a Network General Sniffer for testing, you may simply attach the Sniffer to your test network, and exert fine control on the level of background traffic generated.

The ability of this test to predict real world results are much greater than a single workstation test with no background traffic.

For Measuring Server Throughput

To adequately test server throughput, you will need a minimum of four workstations per server segment to be tested. If you are testing Fast Ethernet, be sure to use 486 or Pentium-based workstations and servers. If you have machines of varying processor speeds, use your fastest system as the server.

You can also use the results of this test to judge a workstation adapter's ability to maintain performance in multi-client environments. In multi-client testing, both Perform3 and NetBench give you individual workstation results, as well as throughput on the server.




Figure 10. Measuring Server Throughput

For Measuring Server Throughput on a Loaded Network

Use the same methodology for generating background traffic in this scenario as you used for generating background traffic in the single workstation test.

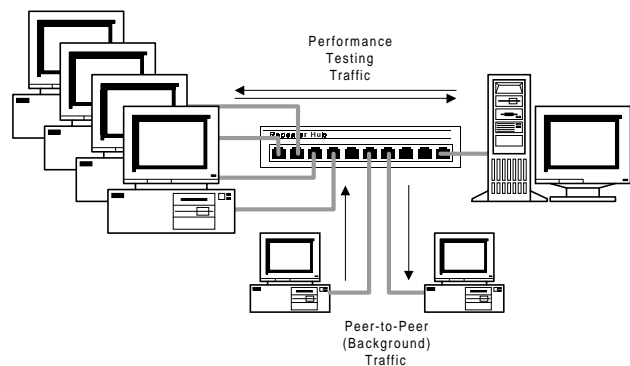


Figure 11. Measuring Server Throughput with Background Traffic

For Measuring Server CPU Utilization

You can measure CPU utilization at the same time you are measuring throughput.

Multi-segment Testing

Many network managers configure servers to service multiple segments, so multi-segment performance testing is important to them. PC Week, in conducting its NIC

Index for server adapters, conducts multi-segment testing, using four segments in its test server, and four workstations on each of the four segments.

You may not have as many test machines as does Ziff-Davis Labs, but you may be able to conduct testing on a two segment network, which can provide some telling data not disclosed in a single segment test.

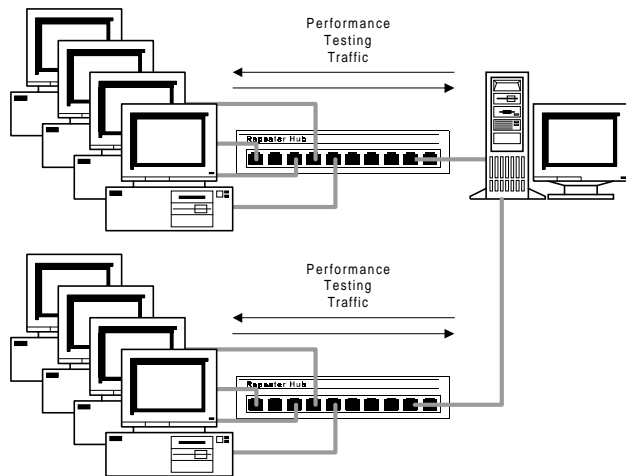


Figure 12. Measuring Server Throughput on Multiple Segments

For Measuring Predictive Pipelining Performance

Predictive Pipelining enables eMASTER+ adapters to achieve 15% better system throughput than other leading adapters on Novell NetWare 3.x and 4.x networks. Predictive Pipelining drivers for servers and workstations are provided on the eMASTER+ distribution diskette as EMxxx_4.LAN and EMxxxP.COM, respectively (where xxx represents either "PCI" or "XSA," depending on whether your adapter is for PCI for for ISA/EISA). To take advantage of Predictive Pipelining, use eMASTER+ adapters along with the Predictive Pipelining drivers in both the client and server computers. For NetWare 3.11 networks, you should add the following line to the server's STARTUP.NCF to enable Predictive Pipelining:

```
SET MAXIMUM PACKET RECEIVE SIZE = 4202
```

Predictive Pipelining results may be generated on NetWare networks using any of the scenarios described above.

About Cogent Data Technologies, Inc.

Since it was founded in 1982, Cogent has become a worldwide leader in ultimate performance LAN products for PCs, including award-winning Ethernet and Fast Ethernet LAN adapters and Fast Ethernet repeater hubs. Cogent products are available worldwide through selected resellers and distributors.

Cogent Data Technologies, Inc.
640 Mullis St.
Friday Harbor, WA 98250
Tel (360) 378-2929
Fax (360) 378-2882
BBS (360) 378-5405
World Wide Web <http://www.cogentdata.com>

Cogent Sales
15375 SE 30th Place, Suite 310
Bellevue, WA 98007
Tel (206) 603-0333
Toll-free in U.S. 1-800-426-4368
Fax (206) 603-9223

Cogent Europe
Rosenstrasse 7
80331 Munich Germany
Sales [49]89/231138-88
Support [49]89/231138-29
Fax [49]89/231138-11

© Copyright 1995 Cogent Data Technologies, Inc. All rights reserved. Printed in U.S.A. Cogent is a registered trademark and eMASTER+ and Predictive Pipelining are trademarks of Cogent Data Technologies, Inc. All other trademarks are the property of their respective owners. Test results shown are specific to a single testing configuration; your results may vary.
Part no. 844-00232-02

Appendix A. Recording Your Evaluation Results

The data below shows performance data generated in Cogent's labs, and presented in the PC Week NIC Index format. This appendix also contains a blank matrix which you can use to record the results of your own benchmark tests.

SERVER NIC TEST RESULTS (NetBench 4.0):
 (AS TESTED BY Cogent Labs, 10/3/95 and 10/4/95)
 ETHERNET CARDS: Netware 3.12 w/ VLM ver.1.20 Rev. A

Vendor/Product (Performance Index Order)	Block Size -->	64	128	256	512	1024	2048	Driver Information		List Price	Larger Is Better	Smaller Is Better
		One Segment, 4 Workstations		One Segment, 4 Workstations		One Segment, 4 Workstations		Name	Date		Perf. Index	Price/Perf. Index
Cogent EM110 TX PCI	Throughput	12.378	21.626	34.530	49.395	62.243	70.880					
	CPU Utilization	12.00%	17.00%	18.00%	31.00%	36.00%	40.00%					
SMC EtherPower	Throughput	10.32	12.72	19.18	15.93	17.29	17.72	EMPCI.LAN	9/13/95	\$189.00	15.53	12.17
	CPU Utilization	12.196	21.170	33.448	46.960	56.874	60.530					
Intel EtherExpress Pro/100	Throughput	11.09	13.23	15.93	16.77	17.23	15.52	SMCPWR.LAN	3/5/95	\$259.00	14.96	17.31
	CPU Utilization	12.254	21.185	33.313	46.456	56.285	60.015					
3Com Corp. EtherLink III (3C595)	Throughput	13.00%	19.00%	27.00%	36.00%	44.00%	46.00%					
	CPU Utilization	9.43	11.15	12.34	12.90	12.79	13.05	E100.LAN	7/5/95	\$169.00	11.94	14.15
	Throughput	12.366	21.561	34.387	49.070	61.405	70.810					
	CPU Utilization	17.00%	27.00%	40.00%	51.00%	61.00%	70.00%					
	Throughput	7.27	7.99	8.60	9.62	10.07	10.12	3C59X.LAN	6/5/95	\$285.00	8.94	31.87
	CPU Utilization											

Throughput presented in Megabits/Sec.

Use the matrix below for recording your own test results.

Vendor/Product	Block Size -->	64	128	256	512	1024	2048	Driver Information		List Price	Larger Is Better	Smaller Is Better
		One Segment, 4 Workstations		One Segment, 4 Workstations		One Segment, 4 Workstations		Name	Date		Perf. Index	Price/Perf. Index
	Throughput											
	CPU Utilization											
	Throughput											
	CPU Utilization											
	Throughput											
	CPU Utilization											

Appendix B. Recording Your Features Comparison

This appendix contains a matrix which you can use to record the results of your features comparison.

Feature	Rating Factor	Cogent	Adapter x	Adapter y	Adapter z
Price*	<i>see below</i>	\$	\$	\$	\$
Automated Configuration	10%	yes			
Diagnostics	10%	yes			
Comprehensive Set of Network Drivers	10%	yes			
Multi-threaded Drivers	10%	yes			
Combo Models	10%	yes			
Multiport Adapters	10%	yes			
Made in U.S.A.	10%	yes			
Performance Enhancement Technologies	10%	yes			
Lifetime Warranty	10%	yes			
Stable Manufacturer, > 10 years experience	10%	yes			
Totals:	100%	100%			

*To factor in price, divide total score
by price (Total/Price):