Network Performance and the Client Connection

A Technology White Paper

May 1995

inside front cover

Introduction

Is workstation performance important? If you buy 486- and Pentium-based PCs rather than 386-based PCs, then you've answered the question yourself - of course performance is important!

As more and more PCs are attached to Local Area Networks (LANs), the performance of the whole LAN becomes as important as the performance of the PC itself, especially as file servers are increasingly being used as a repository for applications software, shared data and mission-critical corporate information. Advances in groupware applications such as electronic mail and Lotus Notes raise the profile of overall LAN performance even higher.

When addressing the issue of LAN performance, LAN administrators typically focus their attention on the network servers, and on the router or bridge connections between LANs. Much effort is often invested in these areas because significant end-user productivity gains can result from improving server LAN connections, and from improving the architecture and performance of inter-network connections. However, in a client-server world, the complementary aspect of network performance, namely **the LAN performance of client PCs** can be just as important as these other factors, yet it is often overlooked or is not considered to be an area where significant improvements are achievable.

This White Paper focuses on the issue of the LAN performance of client PCs, and how this relates to overall network performance and thus to end-user productivity. It demonstrates that using a high-performance LAN adapter to improve the LAN performance of an individual client PC leads to a measurable improvement in that PC's performance.

The conclusion of this White Paper is clear. Using high-performance LAN adapters in client PCs leads directly to high PC performance, and thus to high end-user productivity.

What is Performance?

Three measures of performance are usually used: latency (a measure of waiting time, in milliseconds); frame processing speed (in frames per second); and throughput (in kilobytes per second, KB/s).

Latency is a measure of the time taken to process a packet from input to output of a network device. It is most often quoted in connection with internetworking devices such as bridges, routers and switches. Latency has a strong impact on the throughput of data transfer sessions taking place across an internetworking device, although its effect on performance is dependent on the characteristics of the transport protocol being used.

In theory, latency can be measured for a LAN adapter. However, in practice it is difficult to measure latency across adapters directly, but one can deduce latency by looking at throughput measurements. The impact of high latency in a LAN adapter is usually seen in lower frame processing speeds and lower throughput.

Frame processing speed measures the "tightness" or "efficiency" of the hardware design and the frame processing software of a LAN device. Frame processing speed is typically used as a measure of bridge, router or LAN adapter performance and is measured at small frame sizes to emphasize the effect of the per-frame processing overhead. High performance Token Ring bridges and adapter cards typically process 30,000 packets per second (for 28-byte packets).

Maximum throughput is the figure most often quoted for a LAN adapter. It measures the maximum amount of data that the adapter can transfer per second. This is measured at large frame sizes, typically 1500 bytes on Ethernet and 4 Kbytes on Token Ring. The throughput depends on the Operating System (OS) and Network Operating System (NOS) under test, and for high performance operating systems can approach full bandwidth for a LAN adapter in a file server (1.25 Mbytes/s on Ethernet, 2 Mbyte/s on Token Ring). Performance in a client PC is usually lower than performance in a server.

Several different tools are available to measure throughput, and one of the most popular is Novell's Perform3 utility. This runs under both DOS and OS/2 operating systems, and focuses on stressing the LAN connection.

Server Performance

The performance of a server LAN adapter affects the response times of all the clients that use the server. Consequently, considerable effort is put into optimizing network server performance. This involves using high performance processors, "superserver" PC architectures, fast hard disk drives and high performance LAN adapters. Pipelining technologies referred to in this white paper play a significant role in enhancing server performance. However, this white paper will focus on issues surrounding the client workstations.

Network Performance

Overall network performance is the second area of attention in optimizing network user response times and productivity. This incorporates a number of key elements which are addressed individually below.

Key elements of network performance include:

- Maximum LAN bandwidth
- Bridge / router / switch performance
- Network architecture
- Client LAN adapter performance

Each of the above can be a performance bottleneck, yet each can be optimized to improve overall network performance. Let’s now review each key element.

Maximum LAN bandwidth

A high performance server with a good LAN adapter can serve data at approaching 100% of LAN bandwidth. For Ethernet this is 1.25 Mbyte/s and for Token Ring this is 2 Mbyte/s. Even for a single client using the server this is a limitation; however the throughput of Ethernet or Token Ring is generally judged to be sufficient for today’s applications. For specific high-bandwidth applications, moving to FDDI at 12.5 Mbyte/s may be necessary, but for standard applications the purchase cost and administrative overhead is usually not justified.

On a single LAN segment with many servers, and many clients accessing each server, the total LAN bandwidth may become a bottleneck. Bandwidth utilization on a typical workgroup Token Ring is often around ten percent, indicating that for around 90% of the time no data is being carried on the LAN. When bandwidth utilization figures regularly exceed 50%, the usual solution is to split the LAN into multiple segments.

Bridge / router / switch performance

All bridges, routers and many switches are “store-and-forward” devices, that receive a complete frame into memory before re-transmitting it onto the destination LAN. The additional latency introduced by this store-and-forward operation can severely reduce the throughput of individual sessions across such devices, often by as much as 40% in the case of latency-sensitive protocols such as the very widely used NetWare Core Protocol. If a router must be crossed to reach a server, it is vitally important that the latency of the router is as small as possible. Many product tests and comparisons are published in industry journals and provide data comparing various vendors’ bridges and routers.

Note that published information on the latency of store-and-forward switches cannot be compared directly with that of cut-through switches. This is because the latency quoted for a store-and-forward device is normally quoted as the delay between receiving the LAST byte of an incoming frame and the transmission of the FIRST byte of the outgoing frame, whereas for a cut-through device it is measured from the FIRST byte of the incoming frame to the FIRST byte of the outgoing frame.

For maximum performance, the network needs to be designed to keep traffic on individual segments whenever possible, avoiding hops over store-and-forward devices.

Network architecture

Well-designed networks can produce greater end-user productivity improvements at far lower cost than simply buying larger, faster versions of existing network hardware. Even the fastest router adds more latency than crossing no router at all. To design a good network architecture requires a thorough analysis of the data flows between network clients and servers, taking account of the relative importance of the data traffic and its impact on end-user and corporate productivity. This is a complex topic covered both by many published articles and by specialist network consulting firms.

Client LAN adapter performance

The impact of an individual client workstation on the performance of a real-world shared LAN is often misunderstood. On a large LAN segment, it is natural to assume that with 100 users attached to a 16 Mbps Token Ring, each user has 160 Kbps available to it. In reality, once the user gains access to the ring by capturing the token, data transfer occurs at the full 16 Mbps transfer rate. The limiting factor is the speed at which the Token Ring adapter can transfer data between the PC's memory and the wire.

The following pie chart indicates the various steps taken when a frame of data is sent across a Token Ring LAN from a server to a client PC and the proportions of time spent in each step.

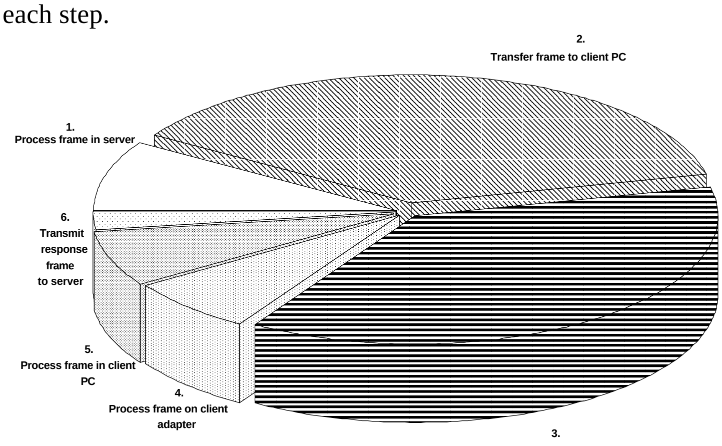

Figure 1: Proportions of time spent in sending a data frame

Single client tests conducted by independent labs confirm the marked difference that the performance of the client's LAN adapter makes to the data transfer speed. Ultimately, this difference lessens the end users' waiting time and increases overall productivity. Steps 2 and 3 in the pie chart indicate that the greatest amount of time is spent receiving the frame and transferring it from the adapter to the client PC. Traditionally, these are performed sequentially, whereas pipelining allows the adapter to transfer a frame to the client PC while still receiving the frame.

We will now flesh out the topic of client LAN adapter performance, and illustrate how influential it is on the real measure of return on LAN investment - increased end-user productivity.

Traffic Profile on a Typical Operational Network

A typical Token Ring workgroup network has a bandwidth utilization of around ten percent. This is compared to a Token Ring backbone which has an average utilization of 50% or more. This does not mean that the PCs on the ring are not taking advantage of the high speed of the Token Ring LAN. In fact, network traffic is very "bursty", characterized by many short bursts of high throughput interspersed with periods of inactivity as illustrated below:

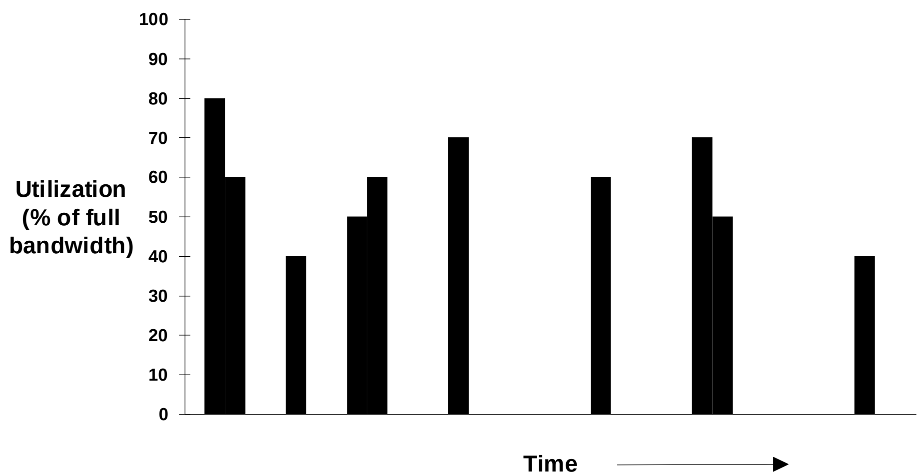

Figure 2: Typical "bursty" traffic on a Token Ring LAN.

With average bandwidth utilization of perhaps ten to twenty percent, a typical workgroup Token Ring is quiet for 80% to 90% of the time. When the ring is quiet, the token simply circles at high speed through the attached PCs, looking for a PC that wants to transmit some data.

When a PC does want to transmit data onto the ring, there is an 80% to 90% chance that the ring is quiet; the PC just waits for the token to arrive which takes a maximum of 17 microseconds on a 100-node ring, or just under one percent of the time taken to transmit a 4K frame. With the PC transmitting onto a quiet ring, the effect is identical to a single-client / single-server environment. This makes performance estimates for a node on a large ring much easier to measure - 80% to 90% of transactions by the client PC will take place in the absence of traffic or interference from the other PCs on the ring. Therefore, most vendors and independent test houses provide comparative figures for their adapters in a single-client environment such as this.

The important conclusion is that performance of a client PC on a large workgroup ring is, for over 80% of transactions, the same as its performance when it is the only client on the ring. Thus, overall network performance is determined to a great extent by the single-client performance of the LAN adapters used in the client PCs.

The Impact of a Fast Client Adapter on Overall Network Performance

Clearly, a fast LAN adapter in a client PC improves the response times seen by the PC's user. Faster response times lead to higher productivity from the user, as he/she is less likely to lose their focus on the task being done. Any excessive waiting time is not only a source of irritation, but increases the chance that the user will be distracted from his/her current task--leave for coffee, etc.

How Can High Client Performance be Achieved?

Many companies buy fast PCs for their users. This is an excellent first step, as internal processing speed and network performance are both improved by using a fast PC processor.

Network performance is improved most dramatically by using a fast LAN adapter in the client PC. Many techniques are used by network vendors to improve the performance of their client LAN adapters, and two of these are described below.

Transmit Pipelining and Receive Pipelining

Two effective techniques used to increase client LAN adapter performance are transmit pipelining and receive pipelining. Both techniques require an in-depth understanding of the inner workings of LAN adapter hardware and software. However, the pay-off is almost a doubling of performance with the consequent valuable gains in end-user productivity. To illustrate how this works, let's look at receive pipelining in more detail.

In a traditional LAN adapter, processes are applied sequentially to an incoming frame. Thus, each process must be completed before the next one can commence, as shown in the timeline below:

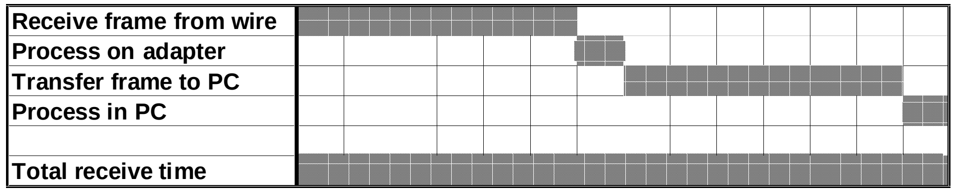

Figure 3: Receiving a frame without pipelining.

Pipelining removes this restriction, allowing processes to run in parallel on the same frame. The two most time-consuming processes, receiving data from the LAN into the adapter and transferring the data to PC memory, can be overlapped to reduce the total time taken to get the frame into the PC's memory. In this way, latency is reduced and performance is increased, as shown below:

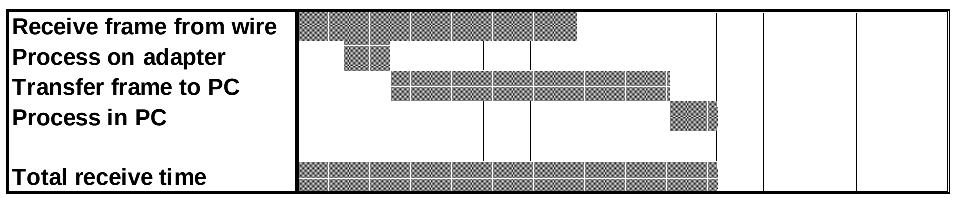

Figure 4: Receiving a frame using pipelining.

Transmitting a frame using pipelining techniques is more complicated, but shows similar performance improvements. Pipelining techniques have been used for several years in Madge Networks' Fastmac drivers for its Smart Ringnodes. Madge has now applied this approach for the first time to shared-memory adapters, with its Blue+ 16/4 ISA adapter. This makes the Blue+ adapter unique in its ability to run with the same driver software as IBM's shared memory Token Ring adapters, and also offer performance improvements of up to 90%. Madge has taken a further step to enhance performance by developing a suite of high performance driver software and its own silicon named RingRunner. This IEEE 802.5 compliant chipset has been uniquely optimized for high performance and is incorporated onto both the Ringswitch Token Ring switch and Blue+ products.

Similar increases in performance are seen in real-life test scenarios. For example, testing by National Software Testing Laboratories (NSTL) using applications such as spreadsheets, database and email on a LAN showed that Madge's Blue+ adapter outperformed IBM's Auto 16/4 ISA adapter. Each test was run with no other traffic on the network and again with background traffic generated to fill 25% of the total network bandwidth.

Spreadsheet Transactions

NSTL's spreadsheet test is a replication of the file system access of Microsoft Excel 4.0. Each call is recreated, including the file, location in the file and size of the operation, in detail, to read and write small (20 KB) and large (100 KB) spreadsheets. This benchmark shows how a network adapter operates under typical spreadsheet usage.

The Madge Blue+ adapter and driver outperforms the IBM Auto 16/4 ISA Token Ring adapter and driver even when the Madge adapter is running with background traffic and the IBM Token Ring adapter is not.

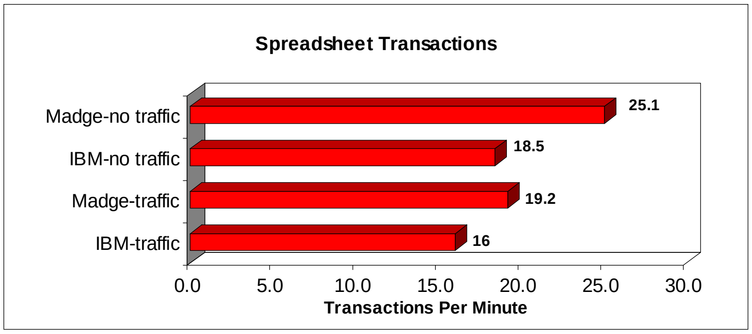

Source: NSTL, May 1995

Database Transactions

NSTL's database test is composed of a Windows database benchmark which reads and writes an indexed database, and displays selected information from the records. The database system stresses the network adapter and driver showing how it functions in a database-intensive environment. The Madge Blue+ adapter outperforms the IBM Token Ring adapter by approximately 15%.

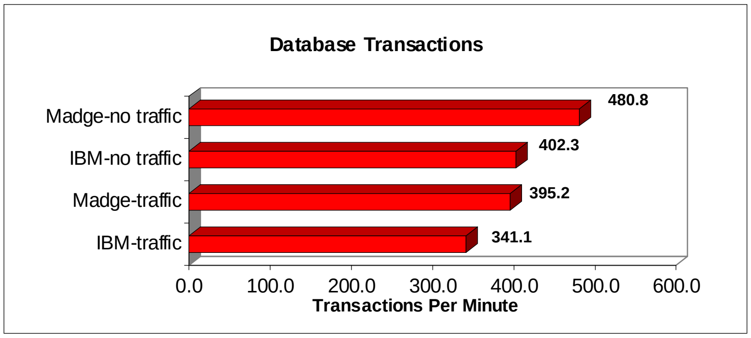

Source: NSTL, May 1995

Cc:Mail for Windows Transactions
NSTL's electronic mail test is comprised of two benchmarks which run under cc:Mail for Windows. Each tests sends a message with an attached file to a specified user. One test attaches a small, representative text file, while the other benchmark attaches a larger spreadsheet.

The Blue+ 16/4 adapter and driver outperforms the IBM Auto 16/4 ISA adapter and driver in both tests with and without generated traffic. For the more common smaller attachments, the performance is even more pronounced as the Madge Blue+ 16/4 adapter performs twice as many transactions as the IBM Token Ring adapter.

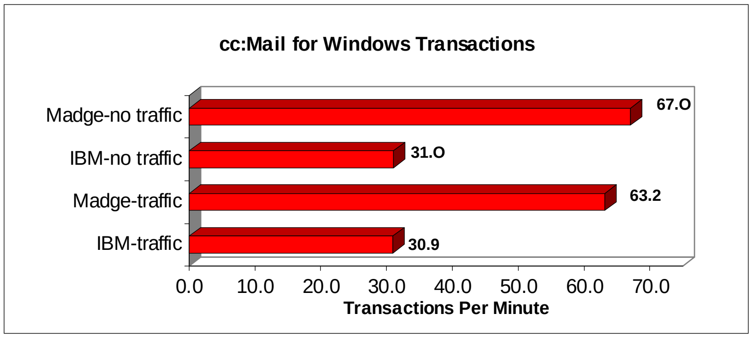

Source: NSTL, May 1995

Above: cc:Mail for Windows Transactions with a small attached file

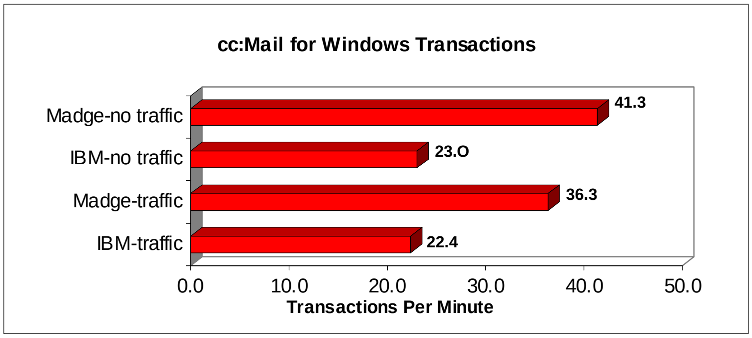

Source: NSTL, May 1995

Above: cc:Mail for Windows Transactions with a large attached file

Summary

Of the many factors affecting end-user productivity on the corporate LAN, the performance of the client PC’s LAN connection is often overlooked. The profile of typical workgroup LAN traffic means that single-station client performance is indicative of the performance that can be achieved in real-life LAN environments. Advances in hardware and software driver technology mean that significant improvements in end-user productivity can be made by using high-performance LAN adapters in client PCs.

LAN adapters using pipelining technology, such as Madge Networks’ Blue+ 16/4 ISA and Madge’s family of Smart Ringnode adapters, do provide real and tangible benefits in typical corporate LAN environments.

inside back cover

Amsterdam
Atlanta
Bangkok
Berlin
Boston
Brussels
Cape Town
Chicago
Cologne
Copenhagen
Dallas
Denver
Detroit
Frankfurt
Hong Kong
Jakarta
Johannesburg
Kuala Lumpur
London
Los Angeles
Madrid
Minneapolis
Munich
Nashville
New York
Oslo
Paris
Philadelphia
San Jose
Seattle
Singapore
Stockholm
Sydney
Tokyo
Toronto
Vancouver
Washington DC

Madge Europe
Madge Networks Ltd
Loudwater
High Wycombe
Bucks HP10 9QZ
England
Tel: +44 1628 858000
Fax: +44 1628 858011

Madge Americas
Madge Networks Inc
2310 North First Street
San Jose
Calif. 95131-1011
United States
Tel: +1 408 955 0700
Fax: +1 408 955 0970

Madge Asia
Madge International Ltd
64-01 Central Plaza
18 Harbour Road
Wanchai
Hong Kong
Tel: +852 2593 9888
Fax: +852 2519 8022

Madge Japan
Madge Japan KK
Believe Mita
43-16 Shiba 3-chome
Minato-ku, Tokyo 105
Japan
Tel: +81 3 5232 3281
Fax: +81 3 5232 3208

Madge Networks reserves the right to change specifications without notice.