Abstract

This document presents a case study that describes the IBM internal NC700 Project which deployed 700 network computers in a token ring wide area network (WAN) to a number of locations in the northeastern United States. The study provides excellent insights into the technical experiences of the IBM team involved in this project. It also shares the results and experiences of the user communities that were affected by the deployment, including the reasons for the project's success. This study is not intended to be a "how to" guide on implementing network computers in a complex WAN environment. It is a case study that is being made available for IBM customers to demonstrate that network computers can do the following:

- Dramatically reduce the costs of ownership
- Replace terminals
- Run personal computer applications, including Lotus Notes
- Be deployed on a large scale in a token ring WAN environment with good performance measurements
- Achieve high user acceptance

IBM customers who review this document and have an interest in exploring possible deployment of network computers in their own environments can obtain additional information by contacting the Document Owner,Ron Wood, at ronwood@us.ibm.com.

Network Computing Introduction

Network Computing is a form of distributed computing that facilitates a new set of network-enabled applications which allow inter-enterprise and intra-enterprise communications, regardless of how those applications are configured or what network transports they use. With Network Computing, the network of computing resources is viewed as the supplier of services. These resources will include resource managers within the enterprise and may extend beyond the enterprise. These resources can be made available from public-domain resources (the Internet) or from commercial providers.

In the past, many of the technology-based solutions have focused on providing users with solutions that help complete tasks faster, better, or cheaper. Graphical user interfaces promised to improve "green screens" and provide better user interfaces, personal computers promised cheaper processing power than mainframes, and distributed processing promised to distribute the workload and provide faster response times. Even though technology has delivered on some of these promises, it has caused several new issues, such as distributed systems management, lower system availability, and endless hardware and software upgrades. As a result, the use of technology as an enabler and as a competitive advantage for business has taken a back seat to the unending quest to catch up with the latest software and hardware releases and the implementation of the latest trend in computing.

It is important for organizations to consider re-engineering both their processes and their strategies, reinventing ways of gaining a competitive advantage. Rapid changes in technology over the past decade have provided tools needed to automate business processes and new ways of doing business that were not available in the past. In the Network Computing model, it will be possible for innovative companies to do things faster, better, and cheaper than their competition, determine how to use technology to change the ground rules of their business, and gain significant competitive advantage in their market segments. The Network Computing model will support the emerging set of customer priorities that include greater speed to market, more flexibility and nimbleness, accelerated global expansion, and more supplier and customer integration. In this model, networks can support interactivity and transport-rich content, which will help customers and businesses redefine concepts such as value, competitiveness, and the very nature of transactions.

The IBM Network Station is based on a true thin client model. The Network Station hardware and firmware are optimized to process and handle graphical user interface events while all other tasks, such as configuration, startup, and management, are allocated to and performed by the server. The network administrators can use the Network Station Manager to personalize each Network Station, provide specific capabilities, and modify these as required from a central location. When the Network Station is started, it downloads the kernel operating system from the server. A user can then log in and the server will authenticate the user's identification and password, allowing the user access to applications.

Even though the cost of the Network Station hardware is cheaper than standalone personal computers, the real cost savings of the Network Station environment are realized in reduced systems management and administration costs and a vastly simplified process for upgrading hardware and software in a production environment. These savings can be directly invested in developing business applications that deliver true business value and a competitive advantage to the organization.

Executive Summary

The NC700 Project involved thousands of hours of work over a five month period by a highly skilled and dedicated technical team from the Global Services' North Geoplex. The project was supported by a staff from the IBM Research T. J. Watson Lab. This work has resulted in the design, deployment, and documentation of a new paradigm in computing with great potential to IBM's Managed Operations business. This summary provides an encapsulation of the major issues and results. Due to the limitations of space, many statements made herein contain facts or findings stated without the associated detailed explanation or context.

Project Overview

The objective of the NC700 Project was to design and evaluate an enterprise-scale application of network computer (NC) technology as a low cost alternative to general office personal computers in commercial managed operations opportunities.

This utilization of NCs as personal computer (PC) replacements offers distinct competitive pricing and service delivery advantages. Initial capital and full life cycle cost per seat appear to be significantly lower than equivalently functional traditional PC desktops. During the NC700 Project and in predecessor work jointly undertaken by IBM Global Services and IBM Research, other significant benefits have been documented by end users, administrators, and organizational management.

NCs are currently regarded by the industry trade press, consultants, and even many NC manufacturers as terminals or single function PC replacement options and not as full function Windows-based PC alternatives. The architecture developed during the NC700 Project proved that this technology can be effectively utilized to deliver reliable, high performance PC functionality in a wide variety of business sites ranging in size from 100 to more than 1,000 local area network (LAN) connected desktops.

Architecture Overview

IBM's first generation Network Station network computer was utilized during the NC700 Project to deliver a Microsoft Windows-based interface to Lotus SmartSuite, Lotus Notes, 3270 terminal emulation, and Internet access for 700 internal IBM office workers across five sites in the North Geoplex. The system architecture consisted of two components: the end-user desktop device and the supporting server cluster. The desktop component included a standard keyboard, mouse, monitor, and NC device. The NC utilized during the project was an IBM Network Station Model 100 that contained an IBM 403 RISC microprocessor, 16 MB of RAM, 2 MB of video RAM, and a token ring port. This desktop device had no internal disk drive or cache, contained no moving parts, consumed only seven watts of power, and was LAN-connected via a standard token ring to applications, files, and computing resources accessed on local clusters of IBM Intel-based servers.

Although local printers can be attached directly to the NC, most of the printers were connected to the network for this project. While other devices may be connected directly to the NC's parallel and serial ports, as well as to its PCMCIA slots, no such devices were connected during this project.

Communication between the desktop device and the server cluster was accomplished via X-protocol on TCP/IP over a token ring LAN. Remote access to this environment can be achieved with modem-equipped PCs via dial-up access utilizing the independent computer architecture (ICA) proprietary protocol. Dial-up performance using this access method is acceptable for casual use, such as remote systems administration, but it is not recommended for full production use.

One server cluster consisting of domain controllers, a file server, and compute servers was installed at each NC700 deployment site (i.e., Southbury, Watson, Poughkeepsie, Rochester, and Somers). The primary design considerations of reliability and flexibility were achieved through the implementation of a redundant, fault-tolerant, and highly scaleable hardware and operating system architecture. This cluster design appears to be well suited, from both economic and performance perspectives, for target installation sites ranging from 100 to more than 1,000 seats. Planned architectural refinements are intended to explore the potential of extending this technology both above and below these ranges.

The operating system, applications, and file system remained on the server cluster. The NC acted as an interface between the display and the shared network infrastructure (i.e., the cluster resources). This centralization of complexity and convergence of standardization are the primary driving forces behind the NC's lower cost potential. The currently available operating system software is a multi-user variant of Microsoft NT 3.51 and provides the end user with a Windows 3.X interface. All applications that run in this operating environment can be utilized, however some modification to single user software may be necessary in an NC multi-user mode.

Java-based technology was not evaluated during this project. As these applications become available, upgrading the existing NC700 desktop device will be required due to the poor performance of Java running directly on the first generation IBM NC. However, minimal changes to the existing server environment will be necessary to accommodate this emerging technology.

Opportunity

Personal computer-based client / server systems that are used in general business environments have extended the reach of computer-based solutions due to their ease of use and the utility of the tools they deliver. Increasing attention is being focused on the true cost of this technology. Full life cycle costs are estimated to be between $6,000 and $12,000 annually. Much of this cost has been associated with the increasing complexity and technical divergence inherent to personal computers and their associated infrastructure.

NC technology offers the opportunity to deliver equivalent PC functionality, performance, and ease of use to LAN-connected clients from shared and controlled resources. This centralization of core computing resources increases the potential for standardization, speeds the deployment of applications, and lowers the support and management costs by offering a convergent technology path without sacrificing the freedom or flexibility demanded by many PC users.

The industry and trade press estimates predict a wide range of total cost of ownership reductions to result from NC technology. Two of the most widely regarded authorities, International Data Corporation (IDC) and Gartner Group, estimate NC savings of 46 percent and 31 percent, respectively, when compared to comparable PC technology. These projections are generally based on minimal cost architectures designed for small installations where the inherent lack of reliability and scalability of a single server are acceptable. Savings for the NC700 architecture are expected to fall within this range of estimates. However, the design of the architecture incorporates significant reliability and scalability advantages necessary for wide scale deployment in managed operations environments. Detailed costing and pricing data is being collected from the NC700 Project as it becomes a steady state installation.

Issues

This new computing paradigm enables a significant migration of complexity and systems management responsibility away from the end user. At the same time, it is possible to provide improvements in functionality, dependability, and performance. Users are no longer responsible for loading and configuring software and operating systems. File systems can be consistently secured and backed up. Policies are easier to implement and enforce. Entire classes of maintenance and support can be significantly reduced or eliminated, and the time and effort to develop and deploy applications can be dramatically lowered.

Critics of this technology have identified several limitations and potential flaws including user resistance to the loss of freedom, network impact, and the lack of a viable NC solution for mobile users. Because applications are loaded on the server and not on the desktop device, end users are restricted to a centrally regulated and controlled suite of software. However, clients continue to be free to configure their user interface and the overall "look and feel."

Logic indicates that replacing a somewhat network-independent PC with a completely LAN-dependent NC would drive up traffic beyond existing capacity. Research indicates that, with files no longer moving across the network, screen image traffic to the NC does not adversely impact well-designed LANs; however, poorly designed networks become obvious with the introduction of large NC populations.

Because NCs are totally network-dependent, a mobile solution is not available today. However, IBM and its industry competitors have agreed on a wireless NC standard and developments can be expected to resolve this issue in the future.

Methodology

This section describes the methodology that was used during the NC700 Project. Descriptions of the various phases, depicted in the diagram, are included.

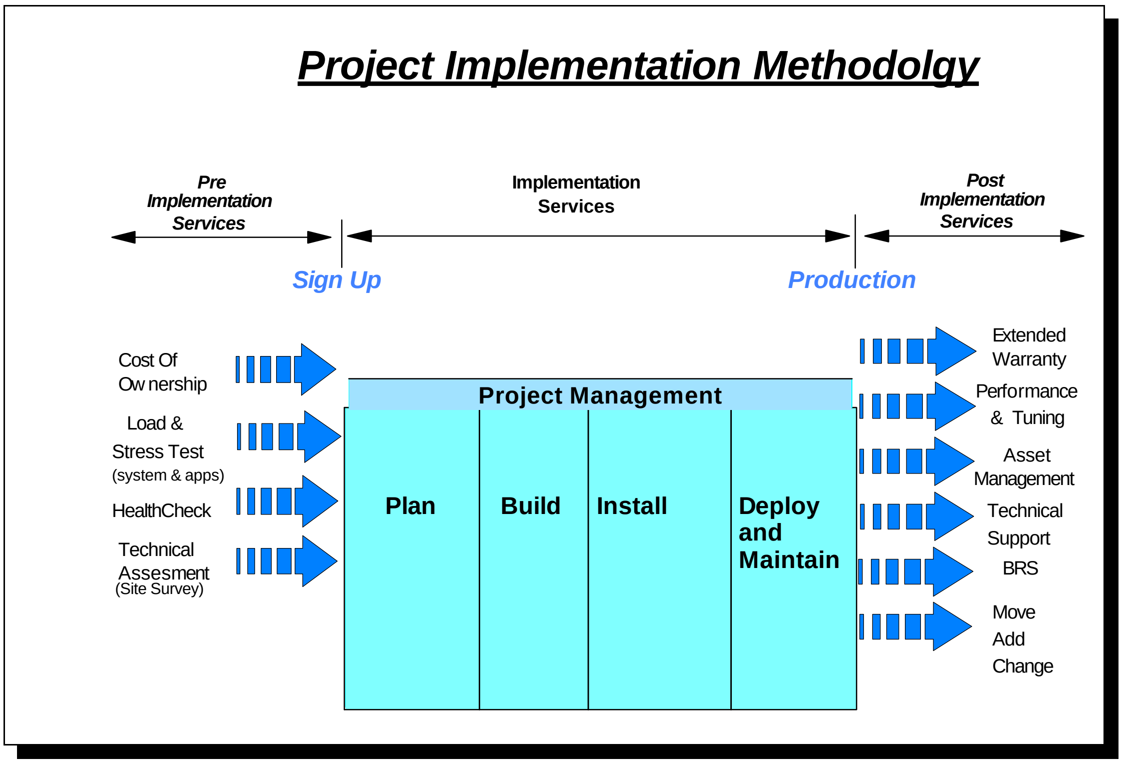

Pre-Implementation Services

Cost of Ownership was used to determine the savings realized over the strategic period of five years. The savings for the NC700 Project were excellent.

Load and Stress Testing was completed to insure that 700 users with varying requirements, including Windows applications, Lotus Notes, and "green screen" usage could be accommodated with performance equal to or better than their current environment. The results in the NC700 project were good.

HealthCheck and Technical Assessment (Site Survey) were done to insure infrastructure readiness for the deployment.

Implementation Services

Plan - An overall project plan was developed to insure successful build and implementation phases.

Build - This phase required technical skills to design, configure, and build an NT WinCenter server environment to accommodate end user requirements of running Windows applications, Lotus Notes, VM, and a Web browser. This phase also included architecting a WAN infrastructure to handle the network traffic, insuring the end user good systems performance, stability, and reliability.

Install - The installation was done with "golden disk" replication methodology. A Solution Build Factory was achieved by replicating each server build and network infrastructure design. This resulted in smooth, seamless installations that met end user requirements and insured a continuous production environment with little disruption of services for the end user.

Deploy and Maintain - The final phase of this project was the deployment of 700 Network Computers throughout the IBM North Geoplex. IBM Managed Operations provides the services, server systems upgrades, and support functions for this Network Computer deployment.

Post Implementation Services

Extended Warranty service was provided on the hardware and software used during this project as it was required.

Performance and Tuning are services that are being provided by IBM Global Services Managed Operations on an ongoing basis to insure a successful deployment.

Asset Management and Technical Support are services provided by IBM Global Services Managed Operations to account for all of the installed hardware, software, networking components, and other equipment required for the successful deployment and maintenance of this project.

Backup Recovery Services (BRS) are provided by IBM Global Services Managed Operations to insure data integrity, disaster recovery, and security.

Move, Add, and Change services are provided by IBM Global Services Managed Operations to accommodate inter-departmental moves, site relocations, and new site installations.

Case Study Introduction

The IBM Network Station is a desktop computer alternative for businesses requiring access to network applications, including those on the host, UNIX, AS/400, Intel, and Macintosh platforms. It promises to lower the total cost of ownership and reduce desktop complexity.

IBM is currently providing the following models of Network Station which are available and can be chosen based on need:

- The Series 100 is optimized to utilize applications running on the server and to allow access to Web browsers. This model is a low cost solution that allows access to multiple servers concurrently and the ability to run Windows applications.
- The Series 300 has all of the capabilities of the Series 100, but it is geared for Internet and Intranet applications. It is optimized for browser performance.
- The Series 1000 has the abilities of the Series 100 and the Series 300, but it was designed for companies wanting to run internally developed Java applications.

For this project, the Series 100 was chosen to allow access to the needed applications, keep costs to a minimum, and prove the concept of the Network Station. WinCenter was chosen as the operating system to be used to allow access to Windows applications. The applications chosen to be used for this pilot included Lotus SmartSuite, Lotus Notes, 3270 access, and Internet access.

This pilot installed Network Stations across five sites in the North Geoplex, which includes the New York area.

Document Content

The purpose of this document is to give the reader information about the results of the pilot test of the Network Stations within IBM. The following topics and information are discussed:

- **System Architecture** - This section describes the components that were used to create the project pilot infrastructure, from the cluster of servers to the end user workstation. It also describes the software needed for the end user.
- **Operating System Description** - This section describes the WinCenter installation that allowed the Network Station to access Windows applications. WinCenter uses WindowsNT 3.51 as the base operating system, which allows the use of 32-bit or 16-bit applications.
- **Application List** - This section lists the applications that were chosen for all of the users of the Network Station for this program and the results of using those applications.

- **Network Architecture and Connectivity** - This section describes several items, including the network protocols that were used in this installation, the ways in which TCP/IP was used and the reasons it was chosen, the use of NetBIOS Name Resolution, and the locations of the cluster servers.
- **Scalability** - This section discusses the number of users that were allocated to each server in order to get optimum performance, the performance of the network and applications, and the measurements of the token rings within the subnet. Graphs and tables that describe howset-up was done and the results that were achieved are included.
- **Performance Issues** - This section discusses the results of using Network Computing Device's (NCD) WinCenter on applications, the results of using the NC with X-Windows, the limitations of memory on servers, and the shortage of hard drive space that can be an issue when people save many files. Improvements that are needed are also reviewed.
- **Cluster Build Overview** - This section covers how the servers were built to support the NC. This Build Process describes, in detail, the building of the Primary Domain Controller, Backup Domain Controller, File Server, Compute Server, and Network Computer.
- **Support Summary** - This section lists a table that shows the personnel needed to make the project like this successful and support the use of the Network Computer. The table includes the responsibilities and required skills for these people.
- **Transition Summary** - This section describes the plans that were created and followed to allow the transition to the Network Computer.
- **Network Impact** - This section includes statistics regarding the performance of the network as it affects both clients and servers, descriptions of the network topology in the area that supported this project, and changes in the LAN traffic before and after the implementation of the Network Computers.
- **Systems Management with Tivoli** - This section reviews the use of Tivoli as a systems management tool for this project, including its use for server monitoring, server administration, scalability (i.e., the ability to use UNIX, Intel, and Alpha operating systems), alert management, and trapping of events and routing to a centrally managed server.

System Architecture

This section reviews the design points and considerations that were involved with the system architecture for the NC700 Project.

Reliability

Over 400 people use NC700 servers every day as their only computing resource. Early problems associated with achieving full functionality of applications, such as Lotus SmartSuite 97 and Notes, have been overcome and, overall, users are very pleased with their NCs. Unlike traditional PC user data, NC700 user data is backed-up automatically on a nightly basis, and data can be restored quickly at the user's request. NCs and individual compute servers are identical and can be replaced quickly without significant configuration consideration. Load balanced groups of compute servers can be set up as a single "virtual server," making individual servers almost like the individual disks in a RAID set.

Scalability

NC700 compute servers can generally support 50 percent more users than originally specified. NC700 clusters can add compute servers to increase the capacity to allow thousands of users to be supported.

Affordability

By creating "virtual servers" that are potentially capable of supporting thousands of users each, the NC700 design makes the most of each server administrator. The NC700 hardware is comparable in cost to traditional PCs. There are potentially substantial savings in end user training, time lost due to standardization, and organized, scheduled roll-outs of new application software.

Network Traffic

NC700 clusters generate approximately 20 percent more total traffic on the LAN than conventional PCs. A large portion of this traffic can likely be contained to a high speed network among the cluster servers. This isolation might leave the general LAN less heavily loaded.

Security

NC700 servers conform closely to IBM security guidelines and to Microsoft's guidelines for achieving Department of Defense C2 security. Users must login in order to access any data. Locked screen savers can be enforced. Software configuration is controlled by user permissions on the operating system and application files and by the WindowsNT registry. The risk of physical theft is minimized as the NC itself is inexpensive to replace and carries no sensitive data with it.

User Desirability

Hundreds of users found the NC700 application list and overall performance good enough to warrant giving up their PCs. In order to make the NC a viable product, thousands of such users are needed and, therefore, the successes need to be expanded.

Application Set

The NC700 Project deployed a full suite of office productivity tools, including an office suite (SmartSuite 97), a Web browser (Navigator 3), e-mail (Notes 4.5), calendar software (Time and Place), and 3270 emulation (PCOMM NT). In order to attract more users, project management, accounting, financial tracking, and other applications need to be included. Given the reliable procedures developed in this project for evaluating and installing applications, deployment of two or even three times as many applications seems like a reasonable goal.

Performance

Even with 50 percent more users than a given compute server was designed to support, users experienced performance that was as good or better than that of the average PC. That performance can be improved through small adjustments to the amount of RAM and disk space on each compute server.

Customizability

Unlike traditional terminal-based architectures, users can customize their operating environment to a great degree with NCD WinCenter. This customization does not require new application software or elaborate user profile maintenance tools, and it does not interfere with the ability to replace individual NCs or compute servers. All user-specific data is kept on a single file server.

Server Architecture Illustrations

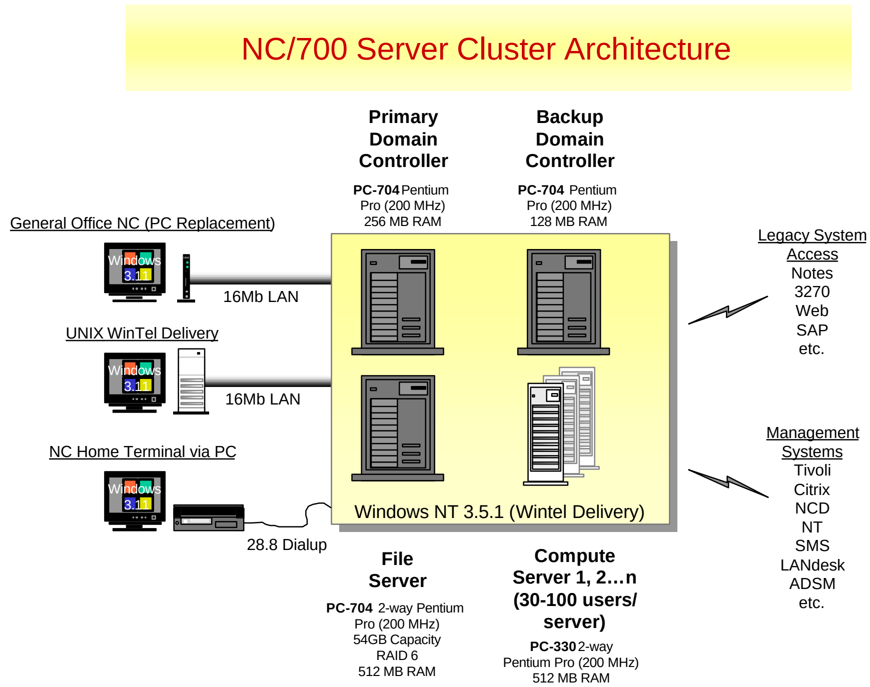

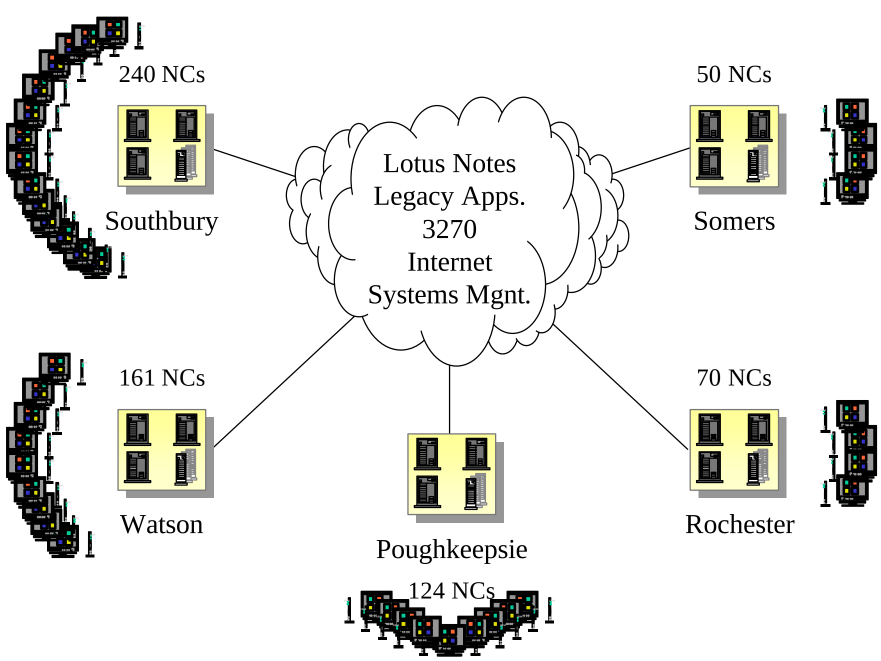

Operating System Description

NCD WinCenter is an authorized extension of Microsoft WindowsNT 3.51. Any hardware or software designed strictly to WindowsNT 3.51 specifications generally runs without issue. However, NT 3.51 is an operating system that is nearly three years old. Fewer new products are being designed for this operating system. By contrast, the Citrix (WinFrame) and NCD, Inc. (WinCenter) extensions to WindowsNT 3.51 that make it WinCenter are young and have some bugs to be worked out. These bugs most often were found in network communications errors and have been addressed on an ongoing basis by NCD, Inc.

Current Issues

The following problems are related to the basic WindowsNT 3.51 architecture and to moving single-user applications to a multi-user environment.

WindowsNT 3.51 Issues

The age of WindowsNT 3.51 brings with it a variety of problems.

- **The Windows 3.x interface is not fully accepted by users.** Once users learn that they will not be responsible for tasks such as configuring hardware and software, this is less of an issue. The old-fashioned window buttons and common dialog boxes definitely make this platform harder to sell.
- **Up-to-date drivers are lacking.** Situations exist where the WindowsNT 3.51 version of a printer driver does not have all of the features of a Windows95 or WindowsNT 4.0 driver. Other common problems are RAID adapter drivers and network drivers. Third party product technical support is also getting thin for NT 3.51.
- **Application support is thin.** When manufacturers are asked for bug fixes, they make it clear that NT 4.0 is their priority. It takes a long time to develop this attitude, even when the problem can be demonstrated to affect NT 4.0 as well as NT 3.51.

Multi-User Environment Issues

Basic WindowsNT is designed with almost all of the components necessary for stable multi-user operation: memory protection, preemptive multitasking, code reuse, and separate storage of each user's application configurations and data. For backward compatibility, however, WindowsNT allows programmers to bypass all of these features. Very few programmers know better than to take the following kinds of shortcuts:

- **Memory Utilization** - 32-bit applications are capable of telling the operating system when critical sections of memory are in use and therefore can load only a single copy into memory, no matter how many copies of the application are launched. 16-bit applications, like the IBM PCOMM terminal emulator, have to be loaded into memory once for every session launched. This does not make much difference to a user who might want to have three, four, or even ten copies loaded, but it makes a huge difference when 15 such users are sharing a WinCenter server. For example, 150 copies of PCOMM running at the same time can take up a total of 450 MB of RAM.

- **Unprotected Memory Space** - 16-bit code does not support memory protection and 16-bit applications can be run in their own memory space, thus protecting other applications and the operating system. However, if 16-bit applications want to communicate with other 16-bit applications, they must share a memory space and can, therefore, cause each other to crash.
- **Multitasking and Processor Capacity** - Rather than being designed to multitask with other applications, many older applications will continuously poll the keyboard and other devices. They will do so whenever given the opportunity, which means that they will consume whatever processor time is available.
- **Configuration Data** - If applications do not make a clear distinction between user-specific and machine-specific configuration parameters, then it is difficult to preserve user settings as a user moves from one compute server to another. This is critical to the redundant design of the NC700 clusters, Microsoft's Zero Administration Initiative, and anything else that requires the encapsulation of user preference settings. Most technical support resources, including those at Microsoft, Citrix, and NCD, will use the term "32-bit application" to indicate that the Windows registry is used appropriately to store configuration data. This is a deceptive use of the term, since 32-bit applications can easily store data in inappropriate locations within the Windows registry and can also write to .INI files. The use of .INI files has been contradictory to Windows programming guidelines for nearly three years, yet it was the default behavior of Visual C++ 4, the premier Windows programming tool, as recently as 1996. It should be noted that the clear delineation of user configuration data and machine configuration data is critical to the deployment of Windows applications on NCs (via WinCenter) and to any application deployment on NCs, including Java.
- **Protection of Executables** - While NTFS, the default file system for WindowsNT, provides access control lists for all file system objects (i.e., files, directories, and links), the operating system and most applications rely on the user write permission in critical application and operating system code directories. The operating system grants such access by default, and applications often fail if it is not available. Unlike UNIX and other traditional multi-user operating systems, Windows assumes that the user has unrestricted access to the system. By default, the operating system only provides protection against other users intruding. On a multi-user system, it is critical that general users do not have write and delete permission for things like the operating system kernel, system settings, and application binaries. Many applications, such as Lotus Freelance Graphics, will not work correctly if the user does not have write and delete permission in the root system directory. Working around these requirements to create a secure working environment is a challenge.

Application List

In general, the installation of applications in a multi-user environment requires distinguishing between user-specific resources and user-general resources. These resources may be user data files (such as templates and dictionaries), configuration data, and executables. Since most Windows applications are not written with this specifically in mind, the designer of a WinCenter environment must often determine these divisions.

The application list for the NC700 Project includes the following:

- Netscape Navigator
- Lotus SmartSuite 97
  - Freelance Graphics
  - Word Pro
  - Lotus 1-2-3
- Time And Place
- PCOMM
- Lotus Notes

This section describes the applications and the issues specific to them.

Netscape Navigator

Netscape Navigator 3.0 is a 32-bit, reentrant application that makes good use of the registry and generally behaves well. Initially, there was some concern that Netscape would consume a lot of network bandwidth when it displayed moving video. This does not appear to be a problem. In addition, the decision must be made between a shared installation of Netscape, which conserves disk space, and individual installations, which permit users to install their own plug-ins. There are three minor outstanding problems:

- Netscape will consume nearly all of a Pentium Pro 200 MHz processor for 3-to-5 seconds when it starts.
- Netscape requires the temporary space for helper applications to be hard coded in the registry, making it impossible to provide secure, personal temporary space on the compute server for this purpose. The NC700 solution used insecure, common temporary space (c:\temp) on the compute server. It is unclear why Netscape does not reference the %temp% system variable each time it requires temporary space.
- Netscape records the addresses for proxy services in hkey_current_user. When a user logs-in to a compute server at a location other than that user's campus, Netscape is temporarily disabled.

Lotus SmartSuite 97

The NC700 Project implemented three applications from Lotus SmartSuite 97. These applications are Freelance Graphics, Word Pro, and Lotus 1-2-3. While all of these are modern, 32-bit applications, they store much of their data in .INI files that are located throughout the SmartSuite directory structure. In all, there are over 40 .INI files, depending on which options are installed. Some of these .INI files hard code the locations of system resources, such as temporary space and the Windows directory. The locations of these resources are changed dynamically by WinCenter and should be referenced by system variables, such as %temp%, or by API calls, such as "Get_Windows_Directory."

Freelance Graphics

Freelance Graphics is also a 32-bit, reentrant application that makes good use of the registry and generally behaves well. There are two major problems:

- In order to display graphics, Freelance needs to be able to use the c: drive for temporary space. This makes securing the c: drive much more difficult. It is unclear why Freelance does not reference the %temp% system variable each time it requires temporary space.
- Freelance will not print to printers defined on remote WindowsNT servers. Therefore, printers need to be defined locally, which requires substantially greater maintenance of compute servers.

Word Pro

Word Pro, again, is a 32-bit, reentrant application that makes good use of the registry and generally behaves well.

Lotus 1-2-3

Lotus 1-2-3 is a 32-bit, reentrant application that makes good use of the registry and generally behaves well.

Time And Place

Time And Place (TAP) is a 16-bit application that stores its configuration data in an .INI file that is located in the services file (%systemroot%\system32\drivers\etc\services) and in the Windows directory. It also creates a c:\CALTEMP directory whenever it is run, making it difficult to secure the c: drive and provide users will secure, private operation.

PCOMM

As noted earlier, PCOMM is a 16-bit, non-reentrant application. It stores its configuration data both in the registry and in a PCSWIN.INI file. Its biggest problem is that it uses a lot of RAM, consuming 3-to-5 MB of RAM per active session. Since some users open as many as 10 sessions, this is a formidable problem.

Lotus Notes

Lotus Notes is a 32-bit application that stores most of its configuration data in .INI files. Notes carries the following challenges:

- Notes expects to find its primary .INI file, NOTES.INI, in the Windows directory, but instead of asking the operating system where the Windows directory is located (using the Get_Windows_Directory API), it hard codes the location in CNODINST.INI. This can lead to the erroneous creation of a shared NOTES.INI file in the system root rather than personal NOTES.INI files in each user's Windows directory.
- Notes allows users to create local copies of databases. If users are allowed to use this function on WinCenter clusters, they can consume a large quantity of disk space.

Network Architecture and Connectivity

Network computers and their servers are completely dependent on efficient, reliable networks. Without them, the NC is nothing but a blank screen. In addition, it is often assumed that network computers, like the early X-stations from which they were derived, consume large amounts of network bandwidth. This makes both the network and the NC vulnerable to abuses by each other. The NC700 Project took several steps to minimize both of these vulnerabilities. During the project, the NC was far more often the victim of unreliable networking than was the network victim of bandwidth-hungry NCs.

In order to provide all of the required functions, NC700 WinCenter servers needed to communicate the following types of information:

- X-Windows - For sending display information to and receiving input from NCs
- Server Message Block (SMB) - For Windows and OS/2 authentication, file and print sharing, and communication among all servers of the project, including five physical locations
- LPR / LPD - For communicating with network printers
- Lotus Notes - For user e-mail and database access
- 3270 Emulation - For access to mainframe applications
- FTP - For general use and migration of user data from legacy systems
- HTTP - For World Wide Web access
- TAP - For Time and Place calendaring
- ADSM - For file server and domain controller backup

To minimize the effect of the NC on the network, and vice versa, efforts were made to make network communication on the NC700 clusters as efficient as possible. This involved the following items:

- Conducting all communication via TCP/IP
- Performing NetBIOS name resolution via name servers instead of broadcasts
- Isolating all intra-cluster traffic on a single token ring switch

The combination of these efforts resulted in a net LAN utilization of only 20 percent more than the PCs replaced by the NCs. Each effort is described below.

Exclusive Use of TCP/IP

TCP/IP was chosen as the only protocol to be used for the NC700 Project because it was more efficient than other protocols and it was the only single transport protocol capable of satisfying all

NC700 requirements. Since each transport protocol includes its own drivers, routing and other issues while using a single protocol are more efficient in the design, normal operation, and troubleshooting phases. TCP/IP's efficiency in large networks is due to its built-in routability, which made it the standard transport protocol of large networks both inside and outside of IBM.

NetBIOS Name Resolution

Server Message Block is the default file sharing and authentication protocol used by WindowsNT and WinCenter and requires the use of the NetBIOS programming interface. The NetBIOS interface is built into the NetBEUI transport protocol, but it can be loaded on top of the TCP/IP and IPX/SPX transport protocols. NetBEUI, however, does not support routing, and both IPX/SPX and NetBEUI require network intensive broadcasting to resolve NetBIOS computer name serving, which is similar to the familiar domain name serving (DNS) used on most TCP/IP networks. Microsoft provides a NetBIOS name server called the Windows Internet Name Services (WINS) with the WindowsNT server. WINS is activated on all NC700 primary and backup domain controllers. Each compute server uses its local primary and backup domain controllers as its primary and backup WINS servers. The primary domain controllers of each site replicate their WINS information with the Watson primary domain controller. Thus, each compute server at each site can resolve the names of all other servers at all other NC700 sites.

Isolation of Intra-Cluster Traffic

Because all user data, including profile data, is kept on the file server, there is a substantial amount of file sharing between the file server and the compute server. This constitutes the bulk of all traffic generated by NCs and their associated servers. Fortunately in the NC700 Project, all cluster servers were located on a single rack and could be isolated on a single token ring switch. None of this SMB traffic affected the local LAN as a whole or the corporate WAN. Unfortunately, other bandwidth intensive systems, such as Notes and ADSM backup servers and printers, were not on the same switch, so this traffic did contribute to WAN and LAN traffic. In future deployments, these services may be segregated differently.

Scalability

The purpose of the measurements was to explore and quantify network computer server architecture scalability and the impact of NCs on network traffic. The goal was to determine the number of users per server that could be sustained and to explore the performance changes and failure modes as the number of users increases. With regard to network performance, the goal was to determine what happens to the network traffic when PCs are replaced by NCs. Aggregate network traffic on the server subnet was determined through measurements on the ring itself.

Server Loading -- Technical Method

Early decisions were made to attempt server scalability measurements using actual users on the production NC cluster in a real operating environment. Simulations frequently give inconclusive results, particularly in predicting worst case scenarios, but it is believed that, even with the uncertainty about how users are actually behaving, more realistic results are obtained from monitoring and measuring performance and resource utilization in a real user environment. Therefore, the strategy was to add users in small increments and observe processor utilization, memory usage, and other parameters using the NT Performance Monitor utility. Also, a measurement suite consisting of timings for various applications, such as cut and paste operations in Word Pro, was utilized. These measurements were done with a stopwatch from an actual desktop NC logged on to the server cluster.

Results

Configuration and Setup

The domain architecture and network setup are shown in the figure below. The machine configuration is shown in the table. At one point in the tests, the disk space was exhausted and an additional 2 GB drive was added to two of the compute servers before continuing the tests. The rest of the configuration remained the same.

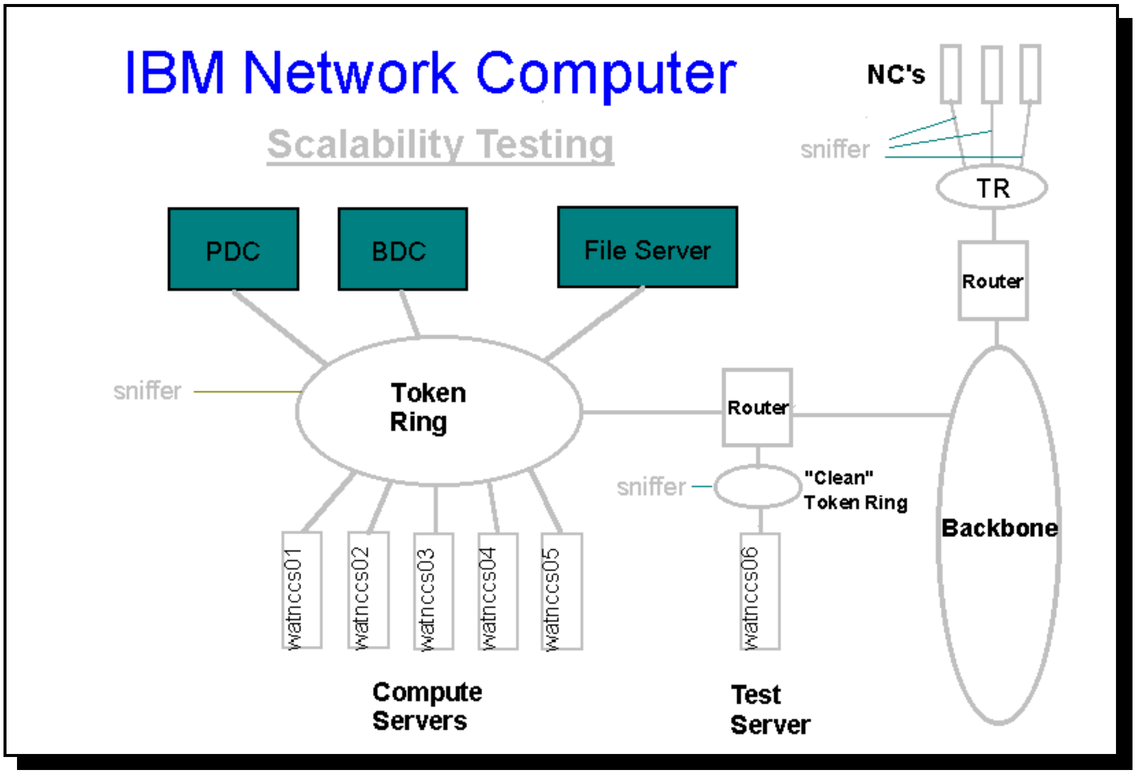

Machine Configurations					
Server	**Model**	**Number of Processors**	**Standard Disk Capacity**	**Model Disk Capacity**	**Memory**
Compute	IBM 330	2	2 GB	4 GB	512 MB
File	IBM 704	2	54 GB	54 GB	256 MB
PDC	IBM 704	1	2 GB	2 GB	512 MB
BDC	IBM 704	1	2 GB	2 GB	512 MB

Processor Utilization, Memory Usage, and Paging

Processor utilization for both processors of the 330 compute server was obtained using the Performance Monitor utility and is illustrated in the following chart. In one set of measurements, the processor utilization was observed at various times during the day on several of the six compute servers. In this way, observations for different numbers of users were obtained. For each point, the processor utilization averaged over 100 seconds.

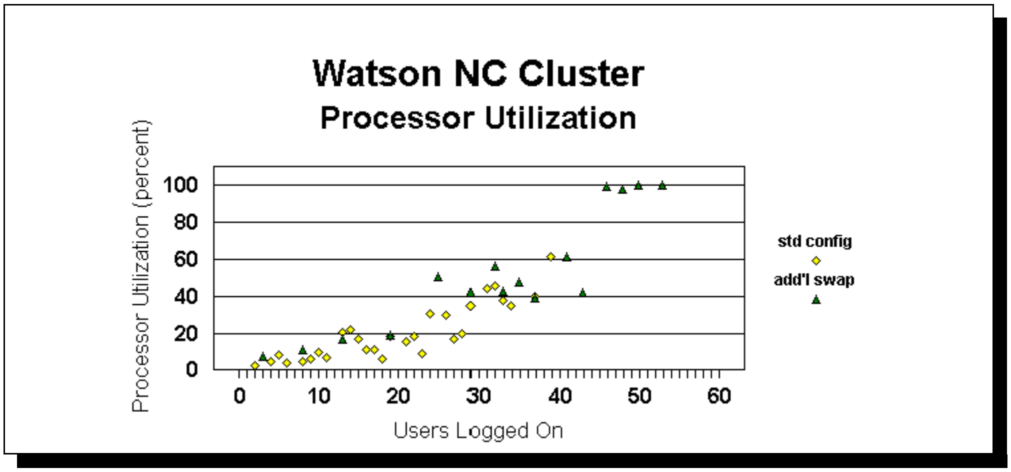

In a different set of measurements, illustrated in the following two charts, Performance Monitor was set to automatically record processor utilization, committed bytes (i.e., the instantaneous virtual memory required and a measure of paging activity), and the number of users as a function of time during the day. Here the processor utilization averaged over 240 seconds. In order to obtain meaningful results by having enough users on a specific processor, the number of compute servers was reduced to two at one point.

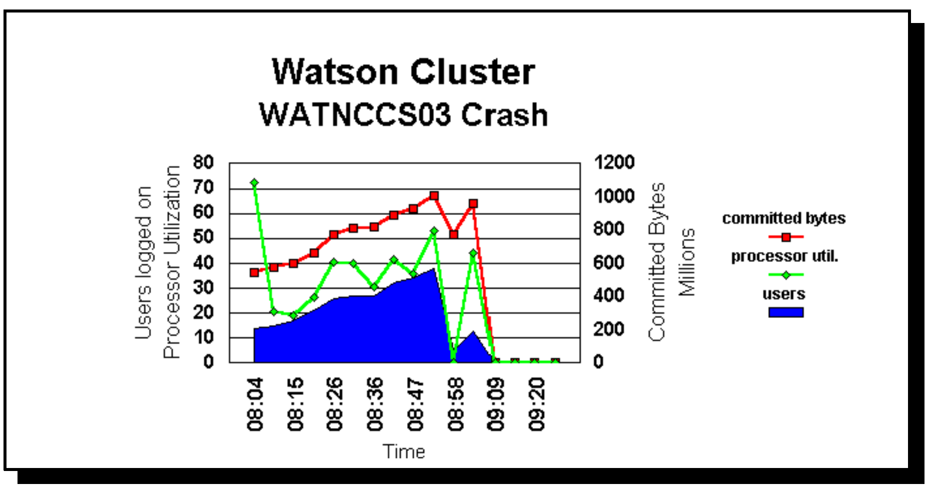

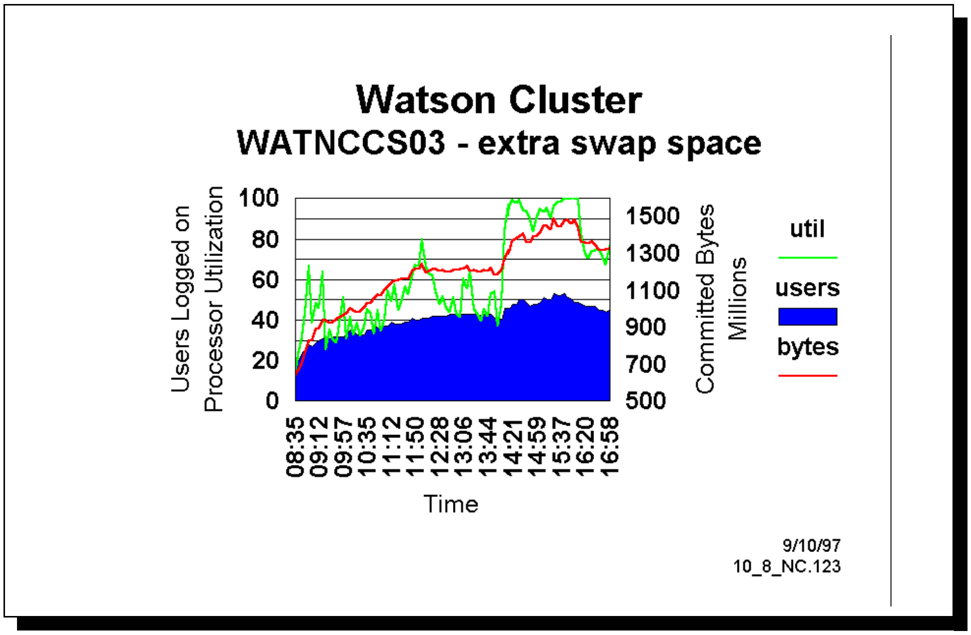

In the chart titled "Watson Cluster - WATNCCS03 Crash" above, the data for the standard and modified configurations is shown. (The differences in the two configurations are shown in the table below.) The composite data with up to about 45 users shows increasing processor utilization to 40-to-60 percent. Ignoring for a moment the 100 percent utilization points beyond 45 users, the projection could be made that about 80 users could be supported, assuming that there were no memory or disk limitations.

Limitations were encountered and resolved. From the standard compute server configuration, data was obtained for up to 39 users. At that point, server crashes were experienced due to insufficient swap space. One incident was captured where the processor utilization, number of users, and committed bytes are plotted as a function of time from about 8:00 to 9:30 a.m. It is believed that server WATNCCS03 ran out of disk space, then crashed at 8:58 a.m. The committed number of bytes exceeded 1 GB, which was all that was allocated and all that was physically available.

After the additional 1 GB disks were added to servers WATNCCS03 and WATNCCS04, the system was entirely stable and data was obtained data for up to 53 users. This number was limited only because all of the available users for that period were logged on. In the chart titled "Watson NC Cluster - Processor Utilization," the processor utilization experienced after the additional swap space was installed jumped to 100 percent when the number of users reached 45. This sudden jump in processor utilization and a rapid increase in the number of committed bytes is seen in other measurements, as well, when the number of users reached approximately 45. These abrupt changes indicate that the system memory is completely consumed, resulting in saturated processors supporting heavy paging activity. It is believed that increased server capacity will result if additional memory is added to the system.

Application Timing Measurements

The effect of the saturated processors is also shown in a series of NC desktop application timing measurements that determined the effect of loading from the user's perspective. Using a stopwatch, measurements were taken for the amount of time taken to open Netscape, Freelance Graphics, and Word Pro. Measurements were also taken for the amount of time taken to bring up a 4.2 MB Lotus 1-2-3 spreadsheet filled with random numbers and a Freelance Graphics chart of moderate complexity. Additionally, measurements were taken for the time to cut and paste a 651 KB Word Pro document as a function of the number of users logged on. As a comparison, the same timing measurements and applications were run on a standalone 133 MHz ThinkPad with Windows95. These timing measurements are shown in following three charts.

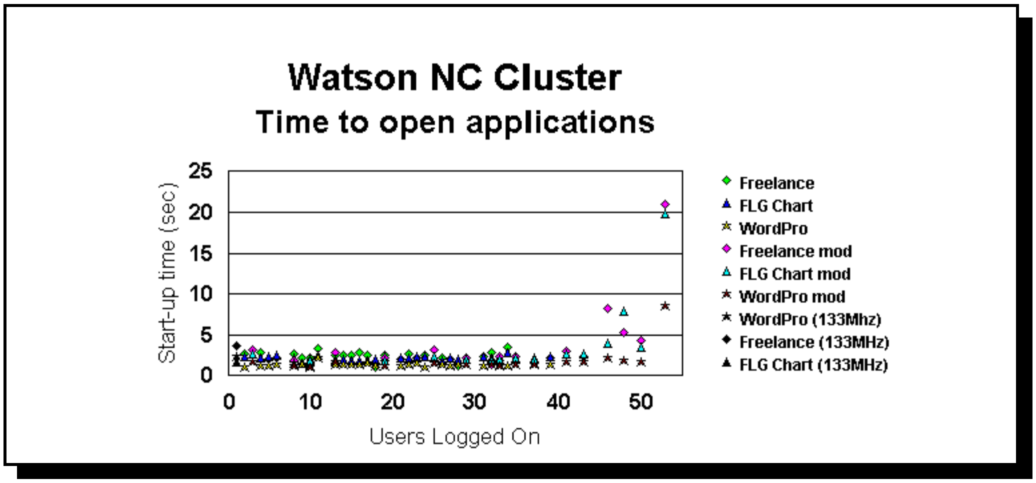

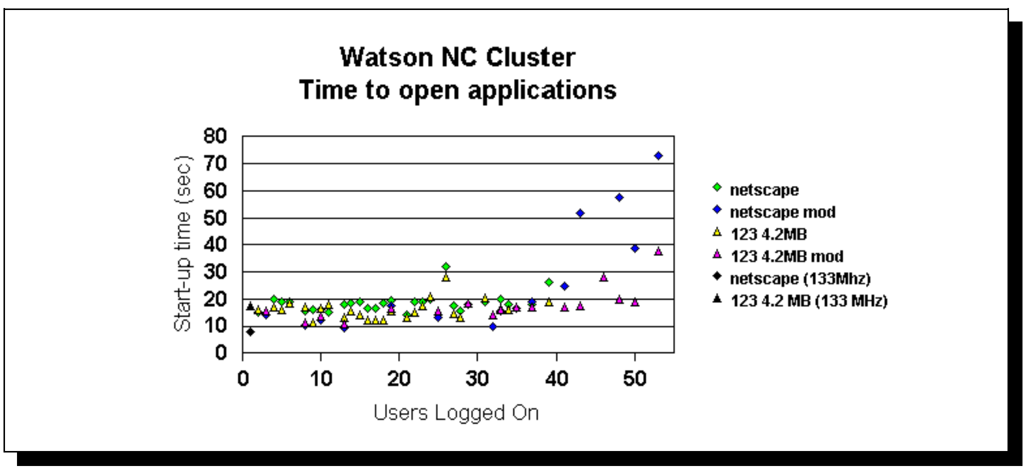

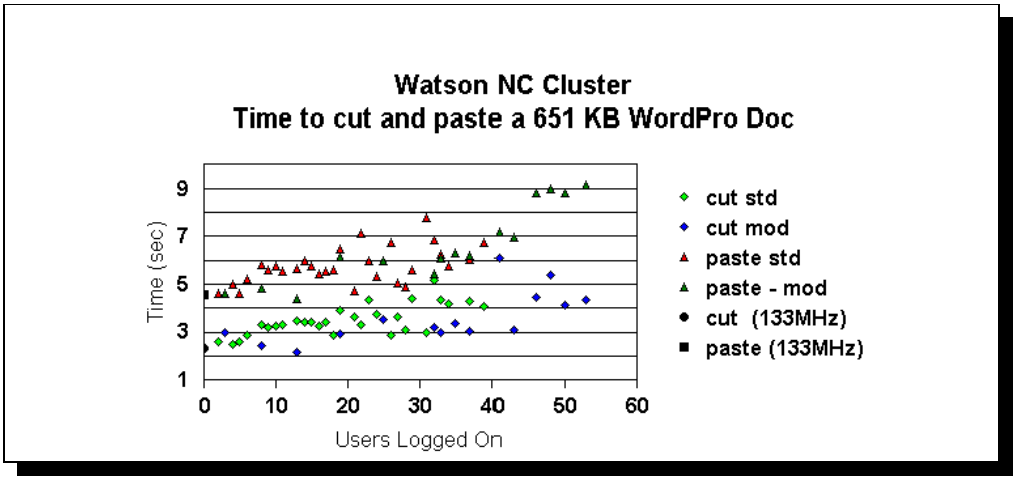

The data is quite consistent with the interpretation of the processor utilization measurements. When the processor is at capacity and has a load of more than 45 users, the performance at the desktop begins to suffer. For example, the time to open the Freelance application, a Freelance chart, and Word Pro are practically constant at less than 3-4 seconds until more than 45 users are logged on, at which time the measurements increase substantially with values up to 20 seconds. The time needed to open Netscape Navigator and a large 4.2 MB spreadsheet also begins to increase and get erratic during times of high user loads. For the measurements of the Word Pro cut and paste exercise, the times generally increase with load, but show a discontinuity at about 45 users.

Measurement Summary and Future Work

These tests probed the capacity of the NC server system in several areas. As originally configured, the dual processor compute servers have enough swap space to support about 35 users. When additional disks are added, this system is stable, but it appears to run out of available memory with approximately 45 users. This is manifested by saturated processors, increased paging activity, and degraded performance, as represented by significantly increased elapsed time needed to open applications and perform operations such as cut and paste. Without the memory limitation, it should be possible to support approximately 80 users per dual 330 processor, assuming that the patterns of usage remain the same.

The plans for further work will involve increasing the memory in the 330 compute servers from 512 MB to 1024 MB. The tests discussed above will be repeated with this configuration. Plans also exist to participate in comparative analysis of four-way compute server technology.

Total Cost of Ownership

Total Cost of Ownership (TCO) analysis was performed for this architecture, with the NC700 Project as input. The objective of this analysis was to determine the cost impact of Network Stations on the total cost of ownership of the standard internal IBM office platform and its associated services. The methodology that was employed included the following steps:

- Determination of the 1998 standard "per seat" rate as a beginning cost
- Identification of the additional costs for the associated services, with focus on the hard costs
- Determination of the current total cost of ownership
- Validation of the current total cost of ownership against the commercial Network Station model
- Implementation of the NC700 Pilot
- Assessment of the cost impact of Network Stations on the current total cost of ownership

Several lessons were learned from the project and the TCO analysis, including the following considerations:

- Architectural Design Phase Considerations
  - Cluster architecture is scaleable.
  - Technology may expose existing network shortcomings.
  - LAN traffic increases approximately 20 percent over conventional personal computers.
- Implementation Considerations
  - Selection of the target audience is critical to the deployment success.
  - Staged rollout is critical, including the use of pilots.
  - WinCenter deployment is not turn-key.
- Financial Considerations
  - Total cost of ownership is the key to savings.
  - Savings projections are in the 13 to 20 percent range.
  - Alternative architectures could drive additional savings.
  - Server infrastructure investment may be required.

The full presentation of these analysis results is documented in Appendix C - Total Cost of Ownership Presentation.

Performance Issues

The NCD WinCenter and the NC700 Project completely overcame two seemingly insurmountable flaws of the NC multi-user Windows NT architecture:

- Refreshing the graphical Windows interface over the LAN should require more bandwidth than the LAN can deliver
- Supporting dozens of users with power-hungry applications should bring even the most powerful PC server to its knees

Processor Utilization

Most users, and even system engineers, do not appreciate that the modern PC spends the vast majority of its time doing next to nothing. The example below shows a 166 MHz Pentium laptop averaging 0.6 percent processor utilization. At that rate, you could have nearly 170 users on a single system before the processor would be fully utilized. This is not literally the case, but the next section will show that dual Pentium Pro systems that were originally designed for 30 concurrent users actually comfortably supported 45 users without modification. With additional RAM, these systems may support over 100 users each.

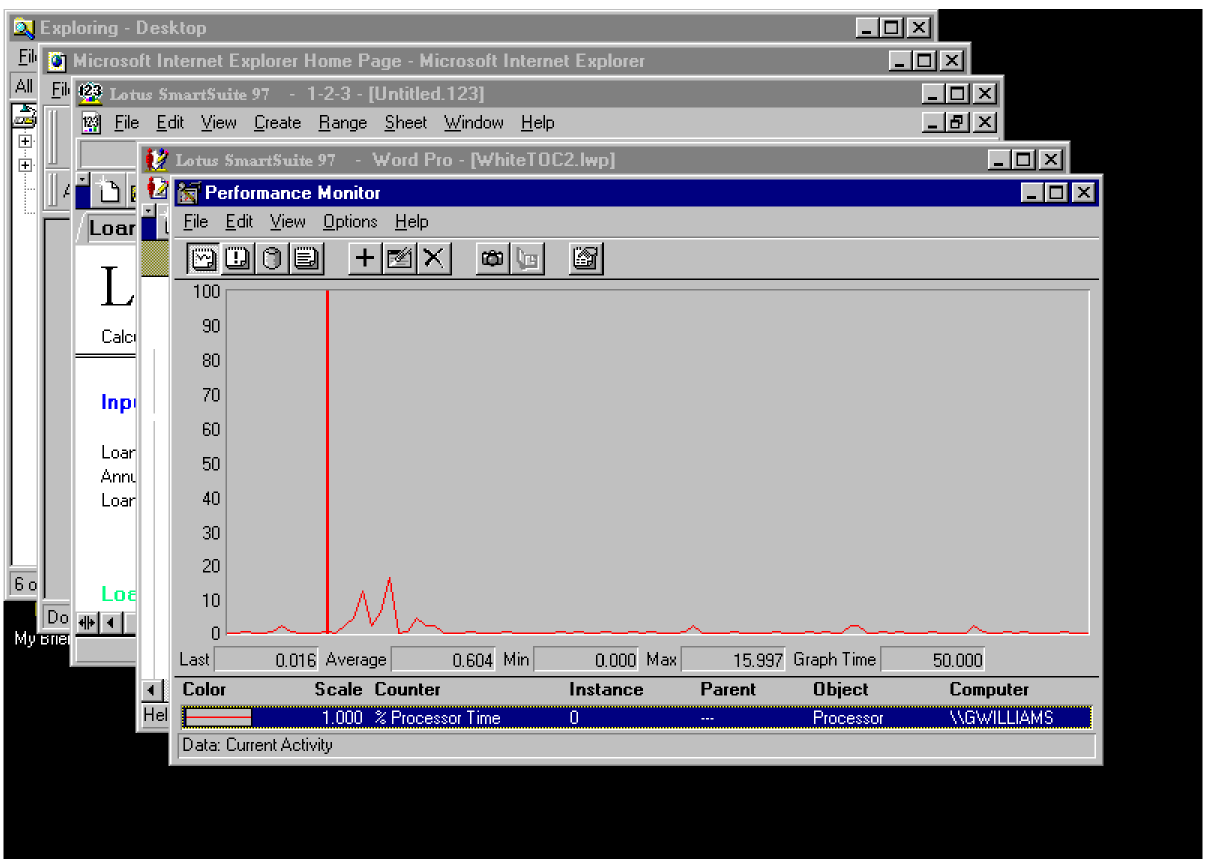

X-Windows Efficiency
Many network engineers remember the tremendous loads X-Windows placed on their networks when it was first introduced in the late 1980s. The impressive item is that X-Windows leaders, such as Network Computing Devices (NCD Inc.), have trimmed this load down to roughly 80 kilobits per second for a 1024 by 768x8 display. It should be noted that measurements taken during the project suggest that the NCD Inc. X-Windows software can be 2-to-3 times as efficient as the software built into the current version of AIX. It should also be noted that file and print sharing traffic consumed far more bandwidth in total than X-Windows did. In most cases, this delivers keyboard and mouse response that is virtually indistinguishable from a normal PC. Display performance degrades as the contents of the display becomes more complex. Overall display performance remains very good to excellent under all observed conditions.

Benefits, Exposures, and Limitations
The overriding conclusion of performance measurements to date is that performance is dependent upon having adequate RAM on the compute server. As physical RAM becomes unavailable and virtual memory is used, performance becomes unacceptable. A more pressing limitation is that of file space. Because WindowsNT does not support directory structures that span file servers, a single home directory structure can only be as large as a single server. This means that cluster systems can only support a finite number of user accounts before the users need to be put into smaller administrative groups.

Memory Limitations
The 32-bit memory model of WindowsNT places an absolute limit of 4 GB of addressable memory per server. While this may seem like a large quantity of RAM, this limit is already being approached; one half gigabyte of RAM was consumed by just 45 users. Doing the calculations, these results lead to 11.4 MB per user and suggests an absolute limit of approximately 350 users on a single system with the 4 GB of addressable memory. The good news is that this is a lot of users. The bad news is that applications will likely grow in size and the 4 GB limit on total RAM is not likely to grow before 1999, when Intel's new 64-bit processor is available.

Disk Limitations
Unlike UNIX network file systems (e.g., NFS, AFS, and DFS) and Novell NetWare, WindowsNT does not support the creation of directory structures that span multiple file servers. If administration is to be simplified by requiring that all home directories fall within a given directory structure, then the number of home directories is limited to those that will fit on one file server. Given that WindowsNT has no means of assigning disk quotas, users are free to fill up available storage as quickly as they wish.

Unfortunately, the highly flexible Distributed Computing Environment (DCE) Distributed File System (DFS) could not be used during this project. Transarc, the manufacturer of DFS systems, does not have a DFS client that is fully integrated with the WindowsNT 3.51 logon procedure. The Lotus Notes 4.5 client does not work at all with the current WindowsNT 3.51 DFS client for unknown reasons. Transarc has expressed an unwillingness to rectify the logon problem, and an improved Notes client will not be available until early next year.

In this project, 200 users filled a 25 GB WindowsNT file server within five months, indicating a usage level of 125 MB per user. Scaleable Serial Architecture disk arrays with over 380 GB of storage can easily handle WindowsNT file servers, suggesting a capacity of 3,000 users per server. However, this is an expensive race, and it is unclear whether hardware developers are faster than users.

These findings suggest that cluster capacities of 1,000 users or more, with over 100 users per compute server, are not unrealistic. The ability to support five times as many users as the largest NC700 cluster, located in Southbury with approximately 200 users, with only a few more machines (i.e., 13 versus 11) should substantially increase the cost effectiveness of future cluster deployments.

Recommendations

The scalability lessons learned in this project can be put to good use immediately. In some cases, such as file serving, the lesson learned calls for both near-term and long-term improvements.

The following improvements can be accomplished in the near-term:

- Increase compute server processing capacity to four processors and upgrade the processors to Pentium IIs as soon as possible (this approach has already been tested under light loads at the IBM Raleigh, NC and Austin, TX sites)
- Increase compute server RAM to 4 GB
- Increase file server disk capacity beyond 100 GB using Scaleable Serial Architecture (SSA) arrays
- Increase compute server network throughput by dedicating one LAN adapter to X traffic and one to file and print sharing traffic (this approach has already been tested under light loads at the IBM Raleigh, NC and Austin, TX sites)
- Increase file server network throughput by increasing the number of LAN adapters and evenly distributing connections to compute servers across those adapters (this approach has already been tested under light loads at the IBM Raleigh, NC and Austin, TX sites)

The following improvement can be accomplished in the long-term:

- Move file serving to DFS

Cluster Build Overview

The build process of the NC700 was, to a large degree, developed and documented simultaneously during the same effort. The actual build of each server component was performed at the IBM Research T. J. Watson Lab, and the team that was responsible for the build consisted of a highly qualified and skilled set of individuals. The process of building the cluster components was documented in a step-by-step procedure.

The NC700 cluster consists of five basic components:

- Primary Domain Controller (PDC)
  - Manages the domain, verifies the logons, and rectifies addressing
  - Able to be a boot server of application server
- Back-up Domain Controller (BDC)
  - Performs the same tasks as the PDC, but it is used as a backup if the PDC is unavailable to the network
- File Server (FS)
  - Contains files needed by the customer applications and files that the customer saves there
- Compute Server (CS)
  - Gives the customer access to the necessary applications
  - May be the WinCenter server to give access to Windows applications
  - Able to be the boot server
- Network Computer (NC)
  - Accesses applications for the customer
  - Consists of the Series 100, 300, or 1000 workstation

The order in which these components are constructed is a definite factor, and they should be built in that order that is listed above.

Build Process -- Primary and Back-up Domain Controllers

These two domain controllers are very similar in every aspect of configuration and physical construction as far as the build process is concerned. There are a few minor differences between them, which are pointed out in the NC700 build documentation. For the purposes of an overview, however, they should be considered the same and redundant. The following list contains the major steps that are required to construct and configure a PDC or BDC:

1. Verify the delivery of all components using the inventory list of the PDC and BDC in the checklist provided with the build section of the NC700 documentation

2. Label the PDC or BDC with the NetBIOS name per the PDC or BDC section of the worksheet as found in the associated appendix of the NC700 documentation
3. Assemble the PDC hardware
   - Install the LANStreamer
   - Install the hard drive
   - Configure and connect the monitor switch box
   - Attach the two power cords to the 704 server
4. Perform the BIOS Levels Check and, when required, perform the updates
   - Note: This step is hardware-dependent and, in the case of the PDC, only the BIOS of the mother board and token ring should be verified.
5. Install the operating system
6. Set up the services
7. Set up the basic directory structure
8. Install WinCenter Connect
9. Install WinFrame Service Pack 5
10. Test the connectivity of the Network Computer

Build Process -- File Server

The process of building the File Server consists of the following general steps:

1. Verify the delivery of all components using the inventory list of the File Server in the checklist provided with the NC700 documentation
2. Label the File Server with the information from the File Server worksheet, found in the appendix of the NC700 documentation
3. Mark the WinCenter Box with the File Server serial number
4. Assemble the File Server hardware
   - Configure and connect the monitor switch box
   - Attach the two power cords to the 704 server
   - Install the hard drives
   - Install the system memory
5. Perform the BIOS Levels Check and, when required, perform the updates
   - Note: This step is hardware-dependent and, in the case of the File Server, the BIOS of the mother board, Mylex adapter, and token ring should be verified.
6. Set up the basic directory structure
7. Install the operating system
8. Set up the services
9. Install WinFrame Service Pack 5
10. Perform the microprocessor upgrade
11. Configure the Mylex adapter
12. Install WinCenter Connect
13. Test the connectivity of the Network Computer

Build Process -- Compute Server

The process of building a Compute Server consists of the following general steps:

1. Verify the delivery of all components using the inventory list of the Compute Server in the checklist provided with the NC700 documentation
2. Label the Compute Server with the information from the pre-installation worksheet included with the NC700 documentation
3. Mark the WinCenter Box with the Computer Server serial numbers
4. Assemble the Compute Server hardware
   - Configure and connect the monitor switch box
   - Attach the single power cord to the 330 server
   - Install the hard drive
   - Install the system memory
5. Perform the BIOS Levels Check and, when required, perform the updates
   - Note: This step is hardware-dependent and, in the case of the Compute Server, the BIOS of the mother board and token ring should be verified.
6. Perform the microprocessor upgrade
7. Set up the basic directory structure
8. Install the operating system
9. Install WinCenter Connect
10. Set up the services
11. Install WinFrame Service Pack 5
12. Test the connectivity of the Network Computer
13. Create a Login Script
14. Create the Install Account
15. Install and configure the applications
16. Duplicate the disk drive for replication
17. Join the domain
18. Inventory all parts that are required for the build of a Compute Server

Recommendations

The current build process can be improved in the following ways, and should be improved for the following reasons.

Increased Number of Applications

The current set of seven applications, while highly functional, does not satisfy enough needs for enough users. The addition of the following applications would dramatically increase the appeal of the offering and increase its potential customer base by at least an order of magnitude. The applications include:

- Microsoft Office
- Microsoft Project

- CA SuperProject
- SAP Client
- ESSBase
- VISIO
- ABC Flowcharter
- Shark

Applications License Metering

It is desirable to have all NC compute server disks be identical, but all users should not be granted access to all applications. Therefore, it would be useful to have a sophisticated license metering facility. Sassafras Software's KeyServer system is already in use in other IBM Global Services managed Windows environments, and it is also in use at several WinFrame installations.

It should be noted that the KeyServer license metering tool will also allow NC clusters to offer users software on a trial basis and enable full access at the stroke of a key. This may offer substantial application sales opportunities.

Support for the IBM WorkPad

For the purposes of this project, there was no mobile solution for NC users. IBM NCs do, however, support the attachment of docking stations for palm-top computing devices. IBM recently signed an OEM agreement with Palm Computing, Inc., and is launching the IBM WorkPad, a new pocket-sized PC companion that uses the PalmPilot technology. Making these devices available as NC peripherals will provide an option for users who currently need laptops.

Common Program Group in Program Manager

Currently, all user program manager icons are defined within personal program groups. When applications are updated, these personal groups must be updated for each user with slow, unwieldy REXX scripts that use Dynamic Data Exchange (DDE) to do search and replace operations in binary registry keys. Moving to common groups for most applications while leaving the users the option of creating their own personal groups allows maintenance of standard program icons to be done centrally and straightforwardly.

With these improvements, the hope is to grow the implementation from the current 700 participating users in the North Geoplex to several thousand, reduce the administrative costs, and increase the quality of the service.

Support Summary

The NC700 Project used four different types of support personnel, described in the following table:

Title	Responsibilities	Necessary Skills
Customer Service Representative (CSR)	Customer service representatives were responsible for all face-to-face contact with the end user. This contact included migration of data from the user's old system to the cluster file server, setup of the NC, initial training, and all subsequent desk-side visits. Since the CSR was usually the first contact a user had with the NC700 Project, and often the only face-to-face contact, their primary requirement was a professional, courteous, efficient bearing.	In addition, the CSR needed to be familiar with: { NC setup { FTP (for user data migration) { Basic NC troubleshooting { Basic WinCenter troubleshooting { Basic TCP/IP troubleshooting { Basic SMB troubleshooting
Help Desk Staff	Help desk staff provided level 1 support for NC-specific issues as well as standard application support.	{ Application features { Basic NC troubleshooting { Basic WinCenter troubleshooting { Basic TCP/IP network troubleshooting { Basic SMB troubleshooting
Site System Administrator	Site system administrators are responsible for identifying any site-specific requirements or problems. Such problems may stem either from the local infrastructure or from the cluster servers themselves. Site infrastructure problems and adjustment of local profiles will be resolved by local site administrators. Solutions to problems with cluster servers will be generated by System Architects and implemented by either system architects or local site administrators.	{ Advanced TCP/IP troubleshooting { Advanced WindowsNT troubleshooting { Advanced SMB troubleshooting { Basic WindowsNT roaming profile configuration { Basic WinCenter troubleshooting { Basic NC troubleshooting
System Architects	System Architects are responsible for the design, build, and third level support of NC cluster systems. Ongoing support includes the creation of updated compute server disks that will be compatible with existing user profiles and file server configuration.	{ Advanced WinCenter design and troubleshooting { Advanced WindowsNT design and troubleshooting { Advanced SMB networking { Advanced TCP/IP networking { Advanced NC configuration and troubleshooting { Advanced application configuration

Skill Transfer -- Deployment to Steady State

At the beginning of this project, there was insufficient expertise available at all levels, largely due to the small installation base of WinCenter and WindowsNT within IBM. Therefore, education had to take place on an informal basis. It is believed that the NC700 System Architect team now represents the best collection of WinCenter experience available anywhere and, therefore should be used for as many deployments as possible. The remaining support resources (i.e., CSRs, Help Desk Staff, and Site Administrator) are site-specific and will have to be cultivated on a deployment by deployment basis.

In order to facilitate skill transfer during the project, site-specific support personnel were brought to the Watson Research Center for several different types of training. Site administrators participated in actual server builds, and all support personnel participated in a two-day NCD training course. Additional skill transfer occurred informally as Watson Help Desk, CSR, and Site Administrator staff interacted with the system architects. Some Watson site support resources, in turn, traveled to the other NC700 locations to transfer their skills to the personnel at those sites.

Recommendations

The NC700 System Architects team is currently being used to create a "Center of Competency" for IBM Network Computing, and formal educational courses are being prepared in Austin. The System Architects team should continue to identify and document procedures and tools to be used by site-specific support personnel.

Transition Summary

The NWS Business Analysis Department requested a Transition Plan based on the NC700 pilot deployed in the North Geoplex. The output of the NC700 Transition Plan will be utilized by the NWS Business Analyst to develop a Network Computer Cost of Ownership Model.

Project Overview

The NC700 Transition Plan Project consisted of two major tasks:

- Collection, manipulation, and verification of data available from the internal roll-out
- Development of the Transition Plan based on the actual data

The intent of the Transition Plan section of the NC700 documentation is to describe the elements of the first task, which includes:

- Data Sources
- Data Manipulation
- Assumptions

This data was then provided to Austin's NWS Transition Support team for the development of the Transition Plan. The data for the NWS Transition Plan was obtained from other NC project plans, Labor Claim Reports, Claim Descriptions, NC Resource Tracking Spreadsheet, and NWS Service Definitions. Refer to the NC700 documentation for further explanation of these items, the plans, and the data.

The original Project Plan upon which the Transition Plans were based was created, revised, and maintained. Although the tasks and hours represent the effort and time spent on the NC700 Project, the start and finish times for some individuals and tasks have been estimated within the range of plus or minus seven days. Although some hourly data came from spreadsheet files, all individuals tracked were listed in the Labor Claim Reports (LCR).

Network Impact

Objective

The network computer model is a variant on the distributed computing model, which is likely to have a different network loading pattern than seen when using full clients. For the purposes of LAN capacity planning and bandwidth prediction on client subnets, it is useful to understand the load presented by individual clients. For the purpose of bandwidth prediction on a backbone or server subnet, it is important to understand the aggregate load on a server subnet. For flat network topologies, understanding the aggregate load would be sufficient.

Method

Two vantage points for data collection were selected: one on the client LAN segment and one on the server LAN segment.

Client

Two series of measurements were taken. To determine the relative difference in client LAN loading, measurement was conducted for five days while the customer used their original workstation, followed by measurement on five additional days using their newly issued NC. The "before" and "after" measurements were not done on consecutive days, due to the complex schedules of the participants. An attempt was made to measure on a "typical use" day. That is, if extremely large or extremely small numbers of bytes were transmitted, compared to other values for that client, the day was excluded. The Sniffer, operating in monitor mode, single station, was used to obtain total byte counts, average frame size, average frame counts, and time. The station filter was defined at the MAC layer and captured all traffic to the measured device, resulting from the tests and other sources. Data was logged every 20 minutes, providing three samples per hour. Byte and frame values represented the quantity seen in the preceding 20 minute interval; utilization and frame size reflected the average value in the preceding 20 minute interval.

Server

Data was collected 24 hours a day, 7 days a week for the duration of the pilot project. The traffic measured on the server segment included client to server (X-Window), compute server to file server, server to server, and transactions with the domain controller components. Traffic due to file backup was also present. In order to be able to measure traffic exclusively related to the pilot project servers, an isolated token ring segment was created, on which was placed five compute servers, two domain controllers (primary and secondary), and one file server. As for the clients, data was logged every 20 minutes, providing three samples per hour. Byte and frame values represented the quantity seen in the preceding 20 minute interval; utilization and frame size reflected the average value in the preceding 20 minute interval.

Watson Network Topology

The following diagram depicts the logical representation of the Watson LAN Topology, with the label "A" indicating the Client measurement points and the label "B" indicating the Server measurement points.

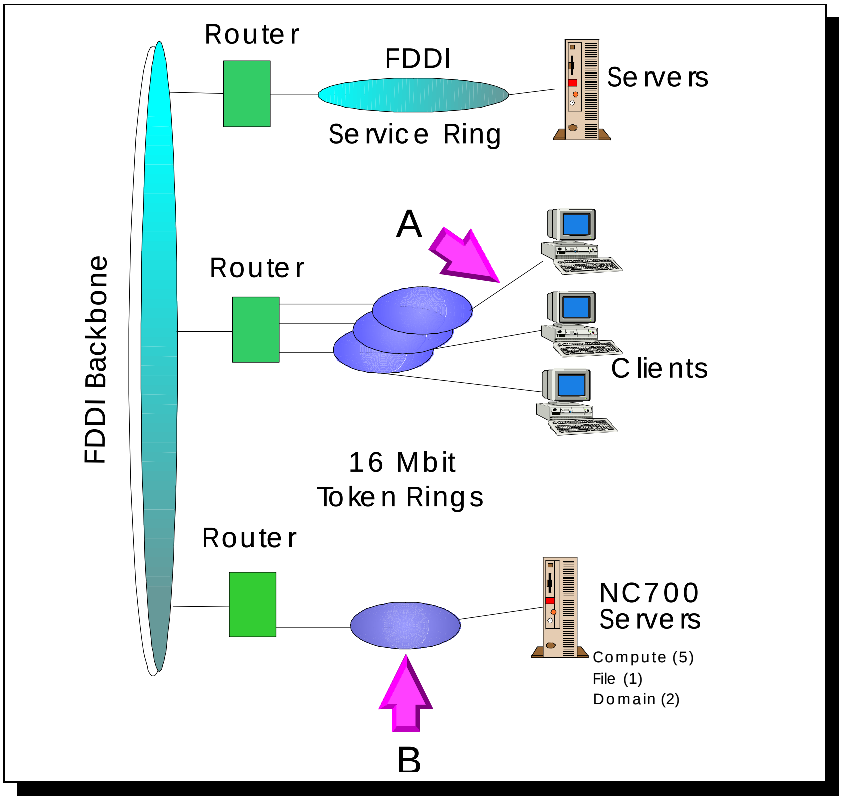

The Watson LAN uses a routed (subnetted), hierarchical topology with an FDDI 100 megabit per second backbone spanning two sites that are ten miles apart (i.e., Hawthorne, NY and Yorktown Heights, NY). Bridge / Routers (Cisco 7000/7500) are interconnected through the backbone and are used to create approximately 80 public IP subnets. The router interfaces connect the hubs (IBM 8260), to which are attached client workstations.

TCP/IP is the major transport protocol, and all applications in this study use IP. There is also a significant NetBIOS traffic component (30 percent) that is bridged using the IP routers. The predominant network interface is the 16 megabit per second token ring. The 10 megabit per second ethernet is also present. While LAN switching is used in some areas, measurements using switched technology were excluded from this study. All clients and servers were attached to the LAN using the 16 megabit token ring.

Results

Client "Before NC" and "After NC" Measurements

The Client data is summarized in the two tables below. The first table presents the average number of bytes during prime and non-prime shifts, the peak number of bytes, and the average percentage utilization of the segment, all based on data for five days on the existing desktop system and five days using the NC. The second table presents the change in load based on the average number of bytes from original to NC client, expressed as a percentage.

Client-Induced LAN Load (in Bytes per Second) and Network Utilization								
Client	**Before**	**After**	**Before**	**After**	**Before**	**After**	**Before**	**After**
User 1	1,181	202	1,042	195	6,145	1,624	0.06	0.01
User 2	458	460	171	7	3,692	3,278	0.02	0.01
User 3	374	2,454	210	2,736	3,588	14,129	0.02	0.13
User 4	1,705	696	625	211	16,721	2,733	0.06	0.02
User 5	276	1,032	45	107	7,673	5,845	0.01	0.03
User 6	424	566	0	56	3,504	1,873	0.01	0.02
User 7	277	301	94	121	2,533	2,136	0.01	0.01
User 8	279	7,253	75	388	2,386	42,049	0.01	0.2
AVG	622	1,621	283	478	5,780	9,208	0.02	0.05

Percentage Change in LAN Traffic after the Conversion				
User 1	-82.9	-81.29	-73.57	-0.05
User 2	0.44	-95.91	-11.21	0
User 3	556.15	1,202.86	293.78	0.12
User 4	-59.18	-66.24	-83.66	-0.04
User 5	273.91	137.78	-23.82	0.02
User 6	33.49	0	-46.55	0.01
User 7	8.66	28.72	-15.67	0
User 8	2,499.64	417.33	1,662.32	0.19
AVERAGE	160.61	68.9	59.31	0.03

Server Segment Measurement

LAN traffic on the server segment is presented graphically for three variables: percent utilization, number of bytes, and frame size. The data was compiled in two views, with one showing hourly values averaged to a single daily value by shift and the other showing the hourly values averaged across all days in the measurement period. The following diagram depicts these results.

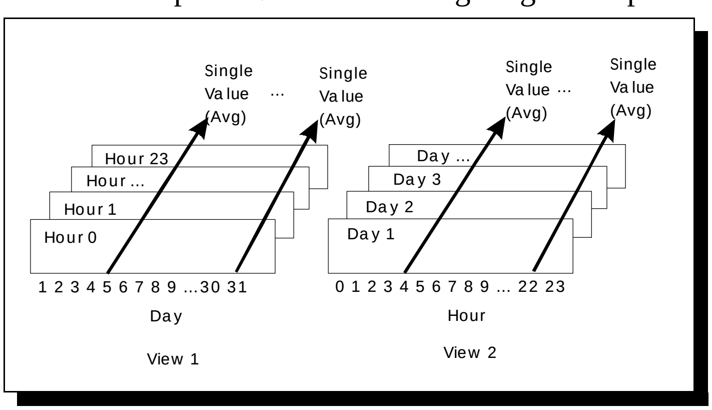

Data is presented in four time periods: June 20 through June 30, 1997; July 1 through July 31, 1997; August 1 through August 28, 1997; and September 2 through September 29, 1997. For the purpose of this short review, focus is on a subset of the data. Data was collected from the server segment during the later part of June 1997 and corresponds to a rather early stage of NC deployment at Watson when the number of users was small. Worth noting are the signature trends of the work week: low utilization of the LAN at the beginning and ending of the week and a maximum in network utilization in the middle of the week. Second and third shift utilizations are rather high and have subsequently been found to result from intense system monitoring.

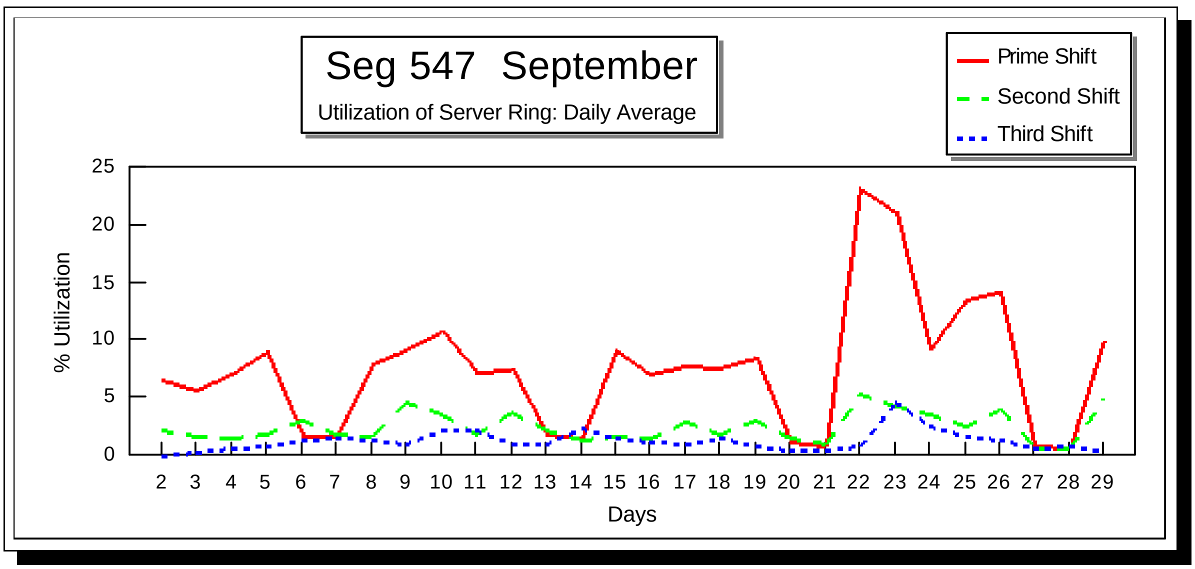

Client

Based on the large variability seen between users and similar large variations between the "Before NC" and "After NC" measurements, it is reasonable to assume that the experimental approach was at least partially flawed. There are several possible explanations to account for the data. Some users will find the new technology novel and interesting, so their use, especially of graphical programs that they may not have previously used, will increase. Other users will be unfamiliar with the new environment and will postpone climbing the learning curve as long as possible. Since the study occurred during prime vacation season, it is also plausible that days off would be followed by unusually high “catch-up” activity, especially in e-mail.

With these comments having been made, the increase observed in LAN traffic generated by the NC during prime shift was 160 percent, during second and third shifts was 68 percent, and during the peak load was 59 percent. The direction of the change was positive for six of eight users during prime shift, positive for five of eight users during second and third shift, and negative for six of eight users during the peak load. The latter anomaly is due to two users who have such large positive values that they reverse the negative values for the other six clients. The peak volume would otherwise have been less for NC clients than for all other starting configurations.

Server

Server network segment utilization is also summarized in the diagrams. There are prominent network utilization spikes that correspond to maximums in weekly usage. While the June data reflects a rather small contingent of NC users, the September data depicts full population of the Watson pilot to about 150 users. An important component of the data is the fact that the maximum average segment utilization for the complete server cluster, serving between 120 and 150 users at peak times, is only in the 25 percent range. Of this traffic, only about 25-to-30 percent consists of X traffic associated with the NC clients. This type of data is quite important when engineering a network to accommodate Network Stations. The balance of the traffic aggregated on the server segment consists of TCP and SMB protocols that are not associated with communication to the NC.

Summary

A series of network measurements has been accomplished on the Watson LAN as part of the NC700 Pilot. These measurements provide planning information for network architects attempting to design facilities for NC client installation. On the server segment measured, typically a 16 megabit per second token ring was sufficient for the installation of the server cluster. The measurements indicate that, for the class of user involved in the pilot, a standard token ring segment can support up to several hundred users if average utilization design data is used. The magnitude of the X traffic allows for subsequent design of those segments supporting NCs only.

Server Architecture Alternatives

The Watson Information Services Department is in the process of reducing to practice two alternative server architectures as they relate to the NC700 Project. At this time, there are production clusters established at various sites. For the sake of this documentation, the T.J. Watson Research facility and the IBM Southbury facility clusters are discussed. The current production cluster architecture is comprised of compute servers, file servers, and print servers with NCs residing locally on the users' desktops. When a user accesses the production cluster, the connection is made through the local area network.

The following are three alternative scenarios that are being investigated.

- **Access to the Production Cluster over the WAN (Phase 1)**
  - This architecture is of interest because it allows centralization of support service and provision of support to smaller clusters of users than does the base architecture outlined elsewhere in this document. Specifically, centralization of the entire server cluster in the IBM Global Services mega-center complexes is desirable from a cost standpoint. Initially, NCs were deployed at the Southbury location and provided WAN access to the Yorktown server cluster. In this initial phase of testing, performance measurements of the connectivity have been started and have confirmed that there is not a noticeable difference in performance as perceived by

the end user. It is believed that this similarity in performance is due to the low latencies experienced by network traffic across the rather robust Advantis internal network. The next phase of the test process will include performing WAN measurements to determine loading of the WAN using this architecture.

- **Local Compute Server Access and WAN Access to File Servers (Phase 2)**
  - After the WAN measurements are completed for Phase 1, Phase 2 will begin, including placement of a compute server (PC 330) on-site at Southbury as part of the Watson server cluster. This architectural test will seek to determine the feasibility of splitting a domain across the WAN. It is conceivable that such an architecture could be used in a mega-center-based deployment. Such a deployment seeks to retain centralized domain control and file service while providing potentially improved application performance to the NC user by placing the compute server locally on the LAN. During this phase, a series of performance and WAN measurements will be taken and compared to the data already compiled.

- **File Server Alternative -DCE/DFS Architecture**
  - The current file server environment utilizes a PC server configured with WinCenter (WindowsNT 3.51) which is formatted using NTFS. Investigation is being made on the concept of migrating from NTFS to a DCE/DFS file server configuration. The benefits of this environment are as follows:
    - -- Elimination of the necessity to manage NTFS data
    - -- Provision of powerful client caching, thereby optimizing performance
    - -- Provision of a single log-on feature to allow users to authenticate over various services and applications
    - -- Administrative ease provided via a centrally managed enterprise file system
  - Note: This alternative file service architecture is currently under construction and evaluation at the T. J.Watson Research Center.

Conclusions

A series of network measurements have been accomplished on the Watson LAN as part of the NC700 Pilot. This study provides aggregate figures which enable planners to size new installations or predict increases when NC devices are added to existing LAN infrastructures. In general, client traffic will be greater than with traditional desktops, so existing design lore will need to be modified and the number of clients per segment adjusted. On the server segment measured, a typical 16 megabit per second token ring is sufficient for the installation of the server cluster. The measurements indicate that, for the class of user involved in the pilot, a standard token ring server segment can support up to several hundred users if the average utilization design data is used.

While the June data reflects a rather small contingent of NC users, the September data depicts full population of the Watson pilot to about 150 users. An important component of the data is the fact that the maximum peak segment utilization for the complete server cluster, serving between 120 and 150 users at peak times, is only in the 25 percent range and the average load is in the 10 percent range. Of this traffic, about 45 percent consists of X traffic associated with the NC clients and another 40-to-45 percent is SMB traffic driven by NC client activity. The balance of the traffic aggregated on the server segment consists of TCP and other protocols which are not material enough to the overall load.

Additional data is required to create a reliable “predictor formula” relating individual client activity to the LAN load. Data is being collected to answer this question.

Systems Management with Tivoli

In the NC700 Project, it became apparent that a tools infrastructure is needed to ensure the NC server clusters are functional or, if not, that an alert structure is put in place to notify key support sources (i.e., Help Desk, LAN Administration, or Third Level). Tivoli Sentry Distributed Monitoring meets this need.

NC Cluster Configuration

In the test scenario, there was an NC cluster that was comprised of a Primary Domain Controller, a File Server, and a Compute Server. The PDC and FS were running WindowsNT and the Compute server was running WinCenter Pro. (Note that the PDC and File Server could also be running WinCenter Pro, which is a derivative of WindowsNT 3.51.)

Tivoli Configuration

The Tivoli server was Intel-based and was running WindowsNT 4. The server was a PDC on its own domain. The Tivoli products that were installed are Framework, Courier, and Sentry.

Policy Regions

Sentry and Courier had respective policy regions. Each policy region had profile managers below it. For instance, in the Sentry policy region's profile manager, there were profiles that dealt with security events and performance monitoring of the NC cluster servers. It can be decided which of the NC servers subscribes to which profiles, created in the Sentry Profile Manager. When certain thresholds were reached, an alert would be logged on the Tivoli server and also sent to the appropriate resources.

Courier profiles contain the configuration data and files for installing software packages in an NC environment. Due to the nature of the NC processing model, the only server that needs to subscribe to this Profile Manager is the File Server, which is the server where the administrator installs the software.

Conclusion

The IBM internal project to deploy 700 Network Computers has clearly demonstrated that we are entering a new age of computing. The 700 IBM users have all learned that, indeed, an inexpensive desktop device can provide the applications they require in order to be productive. At the same time, performance is provided that appears to be occurring locally with easy-to-use graphical user interfaces that provide them with the look and feel to which they are accustomed.

Moreover, the Case Study illustrates that the new world of the Network Computer means users no longer need to be worried about local storage limitations, poor performance, and capacity limitations. The network will, in essence, provide the computing capabilities required to meet the requirements of the end users. Clearly, the IBM Network Stations used in the NC700 Deployment Project show that Network Computers can fill the chasm between dumb terminals and personal computers.

Very key to this report are the following findings:

- The Network Computer Standard Reference definition is viable
- Hardware is less expensive to purchase
- Service, applications, and data are only provided to the Network Station when they are required
- Product life is longer
- Upgrades are less painful
- Reliability is greater
- Software inventory management is significantly reduced
- Standards are easy to enforce
- Each user has a virtual personalized desktop that can be used anywhere in their network environment

We hope that our customers find this report helpful in evaluating the implementations of Network Computers in their own enterprises and reap all of the aforementioned benefits experienced by the IBM personnel who participated in the NC700 Deployment Project.

Appendices

The following appendices are included for reference:

- Appendix A -- Copyright Listing
- Appendix B -- Acronyms
- Appendix C -- Total Cost of Ownership Presentation

Appendix A -- Copyright Listing

IBM, the IBM logo, and Lotus are trademarks of the IBM Corporation. Other brands and products are trademarks of their respective holders:

ABC Flowcharter	Micrografx
AS/400	IBM Corporation
CA SuperProject	Computer Associates
ESSBase	Arbor Software
Freelance Graphics	Lotus Development Corporation
KeyServer	Sassafras Software
Lotus 1-2-3	Lotus Development Corporation
Microsoft Office	Microsoft® Corporation
Microsoft Project	Microsoft® Corporation
Netscape Navigator	Netscape Communications Corporation
NetWare	Novell, Inc.
Network Station	IBM Corporation
Notes	Lotus Development Corporation
NT	Microsoft® Corporation
Pentium	Intel Corporation
Pentium II	Intel Corporation
Pentium Pro	Intel Corporation
SmartSuite	Lotus Development Corporation
ThinkPad	IBM Corporation
UNIX	X/Open Company Limited
VISIO	VISIO Corporation
Windows Internet Name Services	Microsoft® Corporation
Windows95	Microsoft® Corporation
WindowsNT	Microsoft® Corporation
Word Pro	Lotus Development Corporation
WorkPad	IBM Corporation
Zero Administration Initiative	Microsoft® Corporation

Appendix B -- Acronyms

The following acronyms are used in this document and are defined here for clarification.

Acronym	Definition
ADSM	Automated Data Storage Management
AFS	Andrew File System
API	Application Program Interface
BDC	Backup Domain Controller
BIOS	Basic Input Output System
BRS	Backup Recovery Services
CS	Compute Server
CSR	Customer Service Representative
DCE	Distributed Computing Environment
DDE	Dynamic Data Exchange
DFS	Distributed File System
DNS	Domain Name Serving
FDDI	Fiber Distributed Data Interface
FS	File Server
FTP	File Transfer Protocol
GB	Gigabytes
GUI	Graphical User Interface
HTTP	Hypertext Transfer Protocol
IBM	International Business Machines
ICA	Independent Computer Architecture
IDC	International Data Corporation
IP	Internet Protocol
IPX	Internetwork Packet Exchange
KB	Kilobytes
LAN	Local Area Network
LCR	Labor Claim Reports
LPD	Line PrinterDaemon
LPR	Line Printer Client
MAC	Media Access Control
MB	Megabyte

Acronym	Definition
MHz	Megahertz
NC	Network Computer
NCD	Network Computing Devices
NetBEUI	NetBIOS Extended User Interface
NetBIOS	Network Basic Input Output System
NFS	Network File System
NT	New Technology
NTFS	NT File System
NWS	Network Station
OEM	Original Equipment Manufacturer
PC	Personal Computer
PCMCIA	Personal Computer Memory Card International Association
PDC	Primary Domain Controller
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RISC	Reduced Instruction Set Computer
SMB	Server Message Block
SMS	System Management Server
SPX	Sequenced Packet Exchange
SSA	Scaleable Serial Architecture
TAP	Time And Place
TCO	Total Cost of Ownership
TCP	Transmission Control Protocol
TCP/IP	Transmission Control Protocol / Internet Protocol
TR	Token Ring
WAN	Wide Area Network
WINS	Windows Internet Name Server

Appendix C -- Total Cost of Ownership Presentation

This appendix contains the full presentation of the analysis results regarding Total Cost of Ownership for Network Stations.

(Place NSSTOC.PRE here)