SANDIA REPORT
SAND99–0100
Unlimited Release
Printed February 1999

Location Independent Professional Project: A Pilot Study

Marc M. Miller, Jack A. Hudson, and John P. Long

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Prices available from (703) 605-6000
Web site: http://www.ntis.gov/ordering.htm

Available to the public from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

SAND99-0100
Unlimited Release
Printed February 1999

Location Independent Professional Project: A Pilot Study

Marc M. Miller and Jack A. Hudson
Advanced Networking Integration Department

John P. Long
Computer Security Technology Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0806

Abstract

This pilot study project explored the problem of providing access to the nomadic worker who desires to connect a computer through network access points at a number of different locations within the SNL/NM campus as well as outside the campus. The design and prototype development gathered knowledge that may allow a design to be developed that could be extended to a larger number of SNL/NM network drop boxes. The focus was to provide a capability for a worker to access the SNL IRN from a network drop box (e.g. in a conference room) as easily as when accessing the computer network from the office normally used by the worker. Additional study was done on new methods to authenticate the off campus worker, and protect and control access to data.

Acknowledgements

The authors acknowledge the insight, inspiration, and suggestions by Steven A. Gossage and Lawrence F. Tolendino that helped this project to achieve the successes documented in this report.

Contents

Figures

Tables

Executive Summary

This pilot study project explored the challenge of providing access to the nomadic worker who desires to connect a computer through network access points at a number of different locations within the Sandia/NM campus as well as outside the campus.

Three scenarios were envisioned for mobile computing: mobile wired, mobile wireless, and virtual desktop. The mobile wired scenario involves plugging a computer into the network at different locations in a manner transparent to the user. For mobile wireless, transceivers in laptops allow a user to roam much like a cellular telephone. Mobility can also be achieved using Virtual desktop technology by moving the user but not the computer. This technology can make a user's computer desktop appear on another computer.

The bulk of the effort for this project was spent on the mobile wired scenario. Some investigation of wireless and virtual desktop was completed. In addition, authentication issues were investigated.

From market surveys and vendor discussions, it was found that mobile computing using wireless networking is rapidly becoming functional and affordable. Laptops equipped with wireless transceivers can roam buildings and campuses while staying in touch with the network. Data rates have traditionally been low with wireless, typically 1 to 2 Mbps, but the latest technology is reaching 10 Mbps and beyond. However, as data rates increase, signal range tends to decrease. Inter-operability of equipment from different vendors has also traditionally been a problem. This has been mitigated with an industry standard, 802.11, which dictates that equipment built to this specification will operate with equipment from other vendors. Due to security challenges and time constraints, wireless technologies were not implemented for this project. Further investigation of mobile wireless computing is needed. Security solutions appropriate to this technology need time to achieve more maturity.

In addition, virtual desktop technology implementation testing was deferred due to delayed release of Microsoft's Windows NT 5.0.

Investigations for this project show that mobile wired computing can be achieved using dynamic network addressing to assign network addresses as a computer moves transparently around the network. Dynamic Host Control Protocol (DHCP, an Internet Standard Protocol), can be used to provide mobile computing starting in conference rooms and eventually extending to the entire enterprise.

Since DHCP is an Internet Standard Protocol, several commercial products are available. After evaluating some of these products, a product was chosen for implementation in a pilot environment. Several computers in an office area were configured to use DHCP on a daily basis. A primary DHCP server and backup were configured to operate as if in a production situation. Dynamic addressing operated flawlessly, including a demonstration of server redundancy where

the secondary server was able to take over for the primary when the primary was simulated to be inoperative.

Security issues related to DHCP were investigated. The primary issues were identified: who was assigned an address at some particular time, who can make changes to DHCP configurations, and who made changes to DHCP configurations. These issues can be successfully addressed using appropriate DHCP server software.

Using DHCP, network access to a conference room could be greatly simplified. Rather than having to request assistance in preparing a conference room for networking, laptops configured for DHCP could transparently access the network with no need for configuration changes or understanding of networking details.

This concept could be extended to cover the entire enterprise network. The advantage brought about by this approach is reduced effort associated with network address administration. Network address assignment would be handled by the DHCP server. Additionally, DHCP could be configured to respond only to known MAC addresses. This prevents DHCP from interfering with clients that are expecting service from some other source, and provides a mechanism for helping maintain NWIS data by requiring computers to be registered in NWIS before gaining network access.

Authentication issues for the remote mobile user were investigated in order to improve on existing authentication services for the remote user. These issues are user authentication, data protection, and access control. New authentication technologies include PKI (Public Key Infrastructure) and smart cards that can provide strong authentication transparently. Virtual Private Networking, in conjunction with authentication, can provide protected remote access through the Internet, potentially eliminating the need for dedicated circuits. Vendors of security products are now wrapping these technologies into single bundles that can provide one-stop shopping for all authentication needs.

Nomenclature

CCHD	Corporate Computing Help Desk
CIO	Chief Information Officer
COE	Common Operating Environment
COTS	Commercial Off the Shelf (products that are commercially available)
CSU	Computer Support Unit
DHCP	Dynamic Host Configuration Protocol
DCE	Distributed Computing Environment
DSSS	Direct Sequence Spread Spectrum
FHSS	Frequency Hopping Spread Spectrum
FTP	File Transfer Protocol (an Internet protocol that enables file transfers from archives)
HTTP	Hypertext Transfer Protocol
IESG	Internet Engineering Steering Group
IP	Internet Protocol
IR	Infrared (a communications medium using infrared light)
IRN	Internal Restricted Network (Sandia's primary intranet)
ISA	Industry Standard Architecture (an expansion bus used in PCs for plug-in boards such as video and disk controllers and network interfaces)
ISP	Internet Service Provider
Kbps	Kilo-bits-per-second (a measure of the speed of data transfer, 1000 bits per second)
LAN	Local Area Network
MAC	Media Access Control (typically the Ethernet address of a computer that is connected to a network)
Mbps	Mega-bits-per-second (a measure of the speed of data transfer, 1000000 bits per second)
PC/SC	PC/SC Workgroup, a vendor organization designed to address limitations in compatibility of the various smart cards and PCs
PIN	Personal Identification Number
PKCS-11	Public-Key Cryptography Standards #11
PKI	Public Key Infrastructure (an authentication system based on the Public/Private Key system)
PCMCIA	Personal Computer Memory Card International Association (develops standards and promotes the worldwide adoption of PC card-based products)
PSTN	Public Switched Telephone Network (the public telephone system)
RF	Radio Frequency (a communications medium using radio frequencies)
SMS	Systems Management Server
SSL	Secure Sockets Layer
VPN	Virtual Private Network (an encrypted communications path normally through the Internet into an intranet)

Intentionally Left Blank

Location Independent Professional Project: A pilot Study

Introduction

A Sandia/NM worker typically accesses the Sandia/NM computer networks from a network outlet located in the worker's office. The decreasing cost of portable computers has resulted in an increasing number of workers who carry their laptop computers from their office environment to meetings in conference rooms and other offices, as well as on the road. Several workers expressed requirements for accessing the Sandia/NM computer networks from locations other than the office. These workers are referred to as "nomadic" workers.

This pilot study is a response to the needs of the emerging population of nomadic workers. The main thrust centered on the campus-bounded nomad (i.e.: inside the fence), with some time spent investigating the special needs of off-campus access including authentication and encryption. The on-campus nomad is faced with established network infrastructure, network security, and administration policies that require a computer to be at a fixed location, opposing the goal to be mobile. The off-campus nomad has long had the ability to connect to the Internal Restricted Network (IRN) through dial-up, but this area is changing due to the desire to enter the IRN from the Internet (with the need for encryption) as well as new token technology that surpasses the capabilities of the current authentication service.

The pilot study of the campus nomad concentrated on network access from a limited set of sites[a]. The focus was to provide a capability for a worker to access the SNL IRN from a conference room or from the office normally used by the worker with equal ease.

Two other project targets were considered but deferred:

- access to the actual office computer configuration used by the worker was deferred until Windows NT 5.0 becomes available, since meeting this requirement depends on enhanced capabilities to be provided in Windows NT 5.0.
- roaming within a building or a larger outside area through wireless technology such as IR or RF. The study of this capability was deferred until security software matures. Thus a future project may attempt to explore the technology for wireless LAN's (e.g. RF and IR) and to survey some industrial and government installations to

[a] For this pilot project, the campus was defined to be the Advanced Networking Integration Department office and lab, the conference rooms X10 and C2A in building 880, and the CIO office area in building 892; however, one project goal was to examine technology that would be scaleable to the entire SLN/NM campus.

investigate successes, failures, handling of security issues, etc. for a worker who is roaming around a building or area while connected to the local computer network.

For the off-campus nomad, the pilot studied new techniques for authentication, encryption, and access control. New authentication technology includes smart cards that can store a user's authentication data and transparently authenticate. Virtual Private Network (VPN) technology is emerging, allowing a connection to an intranet from the Internet while encrypting the transmission to prevent Internet eavesdropping.

The prototype solution to provide mobility should be scaleable to the entire Sandia/NM campus, be reliable, and be affordable, and be a solution that can be extended to include the off campus nomadic workers. The notes and the subsequent report of the project will describe the actual prototype and the pros and cons of the decisions that were made during the project.

User Input

Network users were questioned about their experiences, concerns, and ideas for the individual roaming around the technical areas and needing access to SNL computer-based information. The project team recognizes that many, many customers have ideas; however, this project represents a pilot study and the scope of the project must be limited. Based on recommendations from management, two individuals were used as examples of nomadic workers with a wide range of experiences and ideas. Ideas and recommendations from management and a number of other individuals were incorporated into the goals based on informal conversations.

The essential points of the user interviews are summarized here:

Currently, network access in a conference room is difficult to obtain and difficult to configure. A typical user should not be required to know of networking parameters such as IP addresses and gateways when attempting to access the network within a conference room.

An ideal conference room environment would consider all aspects of the problems encountered by the nomadic worker. The conference room would provide everything necessary for the network user including power, network connection, standardized plug interfaces, connection to in-room presentation/projection facilities, and a convenient room arrangement. Computers could even be made a permanent part of the conference room, providing simple access to the network. If a user wants more capability than provided by the simple network access, a laptop could be brought into the room, and plugged directly into the network in a simple, transparent manner.

An interesting device was discovered in a public area of a Sandia building that consists of a table that has provisions for power and data connections. Unfortunately, the table was not connected to either power or data, but the design showed promise as a convenient data access point. The small, round table has a roughly circular object in the middle that contains four power plugs and four network plugs. Up to four laptops could be connected to this table and placed on the table or held in the lap.

The basic message gleaned from discussions with users regarding mobile computing was a desire to have easy access to the network from a conference room. Improved network access could be augmented with additional features such as computers, printers and projection devices that are a part of the conference room and network.

Intentionally Left Blank

Location Independent Networking Issues, Technologies, and Solutions

Various issues and technologies related to nomadic or mobile computing were researched for this project:

Issues

- What security requirements are placed on the nomadic worker?
- What static network services might conflict with a mobile computer?
- What impact does mobility have on record keeping such as NWIS?
- How are conference rooms currently providing network access?

Technologies

- How are portable computers being used for communications?
- What technology can best provide mobility?

Security Requirements

How does the nomadic worker affect computer security requirements?

In particular, what is the security policy for how network outlets are used in conference rooms and the requirements for how the outlets are handled when a conference room is not in use?

"The security policy now says that IRN outlets inside the fence but in conference rooms need no special treatment." [1]

This policy statement indicates that a conference room could be configured for mobile computing without changes to security plans or procedures and that a conference room's network ports can be enabled when not in use, thereby always being available for network access.

What are the security requirements that a Sandia/NM on-campus nomad must comply with when accessing the IRN?

"The security policy is that there are no special requirements if the on-campus nomad plugs into an IRN outlet directly with the computer" [1]

Therefore, a nomadic worker can move about the network, such as from office to office, without any changes to computer security plans or procedures.

Portable Computers Used for Communications

Portable computing is increasing in popularity due to handy new gadgets and increasing desire to access information from any location. Portable computers are getting smaller, faster, and cheaper almost daily. Where once a large and heavy laptop was the only way to compute away from the desk, there are now hand-held PCs (H/PC), personal digital assistants (PDA), and just plain small laptops. However, not only is it becoming more and more desirable to compute away from the office, it is also more and more desirable to communicate digitally away from the office. Faxing, e-mail, web browsing, and file transferring are some of the major communications uses today.

Computer communications can be grouped into two types according to the network access method; direct network access and dial-up access. The ultimate goal is the same in both cases, but the methods and equipment are different. Direct network access typically involves a LAN connection using Ethernet, with high-speed transfer of data (10 Mbps or 100 Mbps). This is the typical network architecture in an office environment. However, step away from the office (building or campus) and network access falls into the realm of dial-up. Toss out the fast access and quick response times and deal with 33.3 Kbps transfer rates (or 56Kbps rates in the near future) and slow responses from large, bulky applications that were never designed to be accessed remotely over slow links.

To gain a better understanding of today's portable computer communications needs, the market was surveyed to see what products are available and what products are popular choices for the "road warrior". In addition, project members attempted to become "nomadic workers."

Hand-held Devices

To gain some experience with the H/PC class of devices, a Hewlett Packard 620LX was employed. The H/PC is a fairly new class of computer in which the Windows operating system is used in as small a box as is practical. Called Windows CE, this OS has the look and feel of its cousins Windows 95/98/NT, but is stripped down in order to work within the constraints of a hand-held device (small screen, limited memory, slower processor). Although stripped down, Windows CE is surprisingly useful. Due to the small device size, a stylus is the method of choice for data input, but a fully functional keyboard is included. A finger will even work with the touch screen, but one must be precise when attempting to activate some very small buttons (a finger nail can come in handy here).

The HP 620LX, shown in Figure 1, is built around the Hitachi 32-bit 75MHz SH3 processor. It comes standard with 16 MB RAM, a 256-color, 640x240 backlit full-color display, 10MB user-upgradable ROM, one RS232 serial interface, one IrDA (115Kbps) infrared port, one Type 2 PC card slot, and one compact flash slot. Battery life is approximately 5.5 hours with standard batteries, 11 hours with extended battery, and approximately half these times with a PCMCIA card installed. The single PCMCIA slot is capable of providing various services such as modems, network cards, and external storage devices.

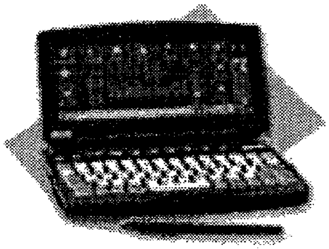

Figure 1. HP 620LX Hand-held PC

Figure 2 shows a typical desktop of an HP620LX. Note that the screen is arranged much like Windows 95/98/NT 4.0, but there just isn't much room. Typical applications are "My Handheld PC" (akin to "My Computer"), "The Internet" (the Internet Explorer web browser), "Inbox" (e-mail), "Calendar" (a scheduler), and "Contacts".

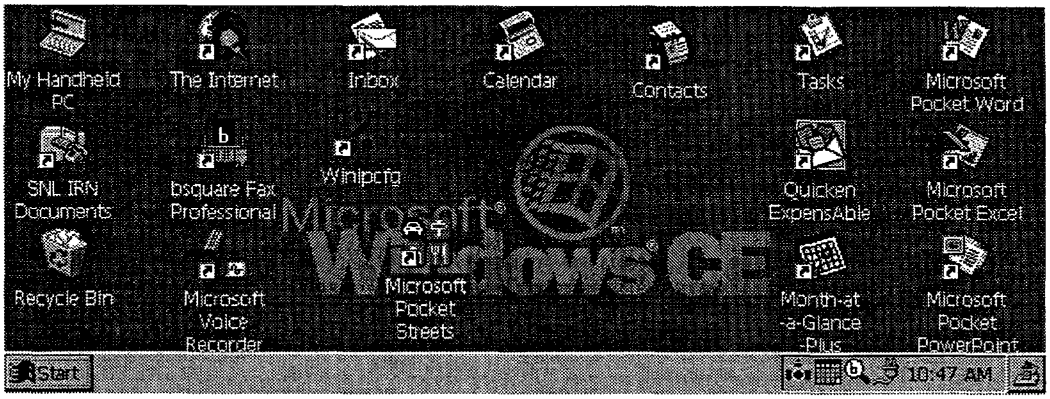

Figure 2. A Windows CE Desktop on an HP 620LX

Synchronization with a Desktop PC

An H/PC running Windows CE can synchronize e-mail, contacts, and files with a desktop PC. To exchange this data, Windows CE Services must be installed on the PC. The H/PC can then be linked to the PC either by a serial port or over the network. This is a handy feature for keeping data up to date in both machines. After travelling with the H/PC and collecting e-mail, simply connect the H/PC to the PC and the two can automatically synchronize; the messages in the H/PC will transfer to the PC where all e-mail is kept. Conversely, when preparing to go on the road, unread e-mail that one might want to read on the plane can be transferred to the H/PC. Replies can then be prepared for delivery when the HP/C is either connected to a network, or synchronized with the office PC.

Dial-Up

An H/PC running Windows CE 2.0 can access remote services through a modem. As of this writing, only a small set of modems is compatible with the OS. For this project, a Hayes Optima 336 V.34 data/fax PCMCIA modem was evaluated, with a street price of $160.

Modem installation is an easy proces with Windows CE, with the process closely following the Windows 95 dial-up networking process.

Networking

The HP 620LX is capable of communicating on a LAN. As of this writing, there is only one PCMICA network card that is certified to function. This is the Socket LP-E, with a street price of $170. The card is designed to use a minimum of power in order to gain as much battery life as possible in the 620LX.

The 620LX is capable of both static and dynamic network address assignment. As with dial-up networking, the network installation process is simple and closely follows the Windows 95 process.

Usability

The HP 620LX H/PC is a useful communications tool. Because many of us are already familiar with Windows 95/NT, this Windows CE device is very intuitive.

Both modem and LAN installations were easy. Connections were easy to establish. Modem transmission speed was typically 33.3 kbps and performed nominally. LAN connections were the standard 10 Mbps Ethernet and also performed nominally.

The dial-up process includes a pop-up window, as does Windows 95/98/NT to facilitate interaction with an authentication server. This makes the 620LX compatible with Sandia's use of the SecurID authentication process. The capability to use both static and dynamic network addressing makes the 620LX compatible with mobile computing as implemented in this project.

The color display, with 640 by 240 resolution, is sharp, but does suffer somewhat with graphics-heavy web pages, where there is not necessarily enough room to show the graphics, and the images are not always clear. The display shows 12 lines.

The keyboard is small, which is a reality for any hand-held device. Rapid typing is impractical for most users. But the full function keyboard does allow most any application to be useful.

The Windows CE 2.0 Pocket Explorer (slimmed down Web browser) has SSL capability for encrypted access to secured Web sites. However, when attempting to access a Sandia secure Web page, such as the telephone directory on the IRN, the message "Unable to establish secure connection" is received and the connection is not successful. This is due to the size of the encryption keys used by the Pocket Explorer. Sandia uses the strongest available encryption and the Pocket Explorer does not. A result of Pocket Explorer's inability to establish a secure connection is that many Sandia Web pages can not be accessed, including time card, telephone directory, and NWIS. This reduces a palmtop's usefulness as a corporate tool when based on Windows CE 2.0.

A potential drawback with the 620LX is the single PCMCIA slot, although this is probably true of most any hand-held device due to the small device size. A single PCMCIA slot results in only a single PCMCIA-type function being available at any one time. For example, after utilizing a modem, it must be removed and replaced with a network card for LAN access (modem/LAN combination cards exist but are not currently compatible with the 620LX). Where this could become a problem with mobile computing is with PCMCIA authentication devices that must operate in conjunction with a modem or network.

Another disadvantage is that although it runs a derivative of the Windows operating system, the 620LX must run programs specifically created for Windows CE.

Laptops

To gain some experience with the laptop class of devices, an IBM Thinkpad 770 was employed (see Figure 3). This is a well-equipped portable computer, built around the Intel 233MHz Pentium with MMX technology. It comes standard with 32 MB RAM (upgradable to 256 MB), a 3.5 in. Floppy Drive, a 5.1GB Hard Drive, a 14.1" display, two Type I/II or one Type III PCMCIA, a built-in sound system, and a 33.6 data/14.4 fax (K56 Flex protocol upgradable) modem. An optional CD-ROM drive is available.

Figure 3. An IBM Thinkpad 770

Laptops have traditionally been running Windows 3.1 or Windows 95, but more and more are running Windows NT or even Linux. For purposes of this project, the Thinkpad was configured to run Windows NT in order to best interact with Sandia's computing resources and security functions. The drawback with NT on a laptop has been a need for a great deal of main memory in the laptop and a lack of power management in NT. Both have been addressed recently with more powerful and better equipped laptops such as the Thinkpad 770, and the addition of power management software to NT.

Compared to a typical H/PC with room for only one PCMCIA slot, a laptop usually has two PCMCIA slots. This provides added flexibility for connecting modems, network cards, external storage devices, or authentication devices such as smart cards.

Dial-Up and Networking

Using a dual-mode 3C562D PCMCIA communications card from 3COM provides both a modem and network interface in one package. This is very convenient for preserving that second PCMCIA slot which might come in very handy for other needs. The 3COM 3C562D has a street price of about $260.

Installation of the 3C562D is simple. The network software must be installed first. Because the laptop was configured to use a PCMCIA card management application, by plugging the 3C562D into the PCMCIA slot, an installation program started automatically, since no driver had yet been installed. Simply following the directions on screen resulted in the proper driver being loaded for the card. The network parameters were then entered. In this case, the network adapter was configured to use dynamic addressing and therefore no further configuration information was needed. The laptop was rebooted and the network was operational.

To configure the modem, it was only necessary to start Dial-Up Networking and follow the directions. The modem was automatically detected. The Dial-Up Networking configuration needs to know what number to dial and what protocols to run and then the service is ready.

Usability

A laptop configured with modem and network adapter operates very much like a desktop PC, and when compared to some older PCs, even better. For example, the Thinkpad used for this project is overall a better machine than the desktop PC used to create this report. As an experiment, the Thinkpad was put in place of the desktop PC and was found to be a good replacement, with the exception of not having enough main memory (32MB rather than 64MB or more, which is easily corrected).

Conference Room Network Access Trials

Several conference rooms have network access available, but it is not a trivial process to take advantage of this conference room feature. If a room feature is too difficult to use, it will not get used.

In an attempt to gain knowledge of how network access in a conference room can be improved, a baseline was established by utilizing the current facilities. The first trial utilized a single conference room with the desire to access network resources such as file servers, and the internal and external Web. The second trial utilized two conference rooms with the desire to access network resources and the user machines in the two conference rooms.

First Conference Room Trial

The conference room in Building 880, room X11 was tested for network access. The room was scheduled using the Web-based conference room scheduler. The details for the room indicated that it was set for network access, with the contact listed as the CCHD. The CCHD was called and network access was requested for the desired date and time. The help desk person didn't know immediately what to do, but called back a few hours later with the answer. After one false start trying to get the answer, she had contacted the 880 CSU with the request, and the CSU was then to provide details of how to utilize the network from the conference room. The CSU said the room was ready, but a few network components would be needed (a small network hub, a transceiver, and cables; necessary because the active drop was on optical fiber and needed conversion to electrical), and these were delivered a few minutes later. The CSU person gave an explanation as to how to connect the equipment and what IP address to use and offered to be at the conference room to help with the configuration, but the offer was declined due to sufficient knowledge of the operation of the equipment. On the scheduled day, the equipment was brought into the room and connected together and the PC configured with the supplied IP address. One piece of information not supplied by the CSU was the router address, but this was assumed to be host number 254, as most routers are given this address. With the configuration complete, the network was successfully accessed and the test was complete. The CSU person came by the next day to pick up the equipment, asked if the network access was successful, and was advised that the network access worked fine.

This test showed that network access in a conference room is possible but not without considerable effort. In addition to scheduling a room, time must be spent interacting with a support organization such as the CSU to arrange for network access. Scheduling a conference room for network access is also totally dependent on the availability of a CSU person.

Second Conference Room Trial

A second conference room network access test was conducted involving two conference rooms spanning two CSUs. The purpose was to see how two different CSUs would operate in this situation and how network access would perform across different network segments.

The first step was to schedule two conference rooms for the same time and date. These conference rooms must have "computer access" as defined in the web-based Conference Room Scheduler application. The conference room in Building 880, room X10 was chosen since it was close by. Another room was found in nearby Building 891, room 3098, and the two rooms were scheduled.

The next step was to request that network access be made available for both conference rooms for the scheduled date. From previous experience with 880/X10, the CCHD was called to arrange for network access. But for 891/3098, an individual was referenced. This individual was called to make the arrangement and she didn't know how to do this. She later called back and said to call the CCHD. The following time line shows the events needed to make the network arrangements:

> 3/17 Conference Room Scheduler application was used for scheduling two conference rooms, 880/X10 and 891/3098, for March 19th at 10:00 A.M.
>
> 3/18 2:54 p.m. Initial contact was made with the manager of conference room 891/3098. She did not know about computer access, needed to find out, and would call right back with the answer.
>
> 3/18 3:04 P.M. The conference room manager called back. She determined that the 891 CSU should be contacted. How does one contact this CSU? There was no information on the internal Web regarding how to contact the 891 CSU. There were only telephone numbers for individuals within the CSU. Rather than randomly calling people to get an answer, and since the CCHD was to be contacted for 880/X10, why not ask them about 891/3098 as well? This was strictly a guess, but an educated one from experience; a customer should not have to guess how to obtain service.
>
> 3/18 3:09 P.M. The CCHD was called and told about the need for computer access to two conference rooms. Two service ticket numbers were given; one for 880 and one for 891. Someone from each of these CSUs would then call.
>
> 3/18 3:27 P.M. A call was received from the 880 CSU ticket dispatcher. She verified the request for computer access to 880/X10 for March 19th at 10:00 A.M. She advised that a CSU technician would be available to assist in setting up a network connection in the conference room.
>
> 3/18 4:05 P.M. A call was received from the 891 CSU ticket dispatcher. He verified the request for computer access to 891/3098 for March 19th at 10:00 A.M. He advised that a CSU technician would call to arrange to meet prior to the meeting for setting up a network connection in the conference room.
>
> 3/18 4:25 P.M. A call was received from an 891 CSU technician. He advised that he wasn't sure if the network drop in 891/3098 was active and would have to check on it and call back. He also said he would be available tomorrow morning to help configure the laptop.
>
> 3/19 8:00 A.M. 891 CSU technician had left a voice mail advising that the network drop in 891/3098 was active and ready. He suggested he be called prior to the meeting and he would be there.

3/19 8:22 A.M. 891 CSU technician called to verify receipt of his voice mail and to advise that when he is to be called, he may be in a meeting but to leave voice mail and he will be paged.

3/19 8:41 A.M. 880 CSU technician called to give an IP address for 880/X10. He said he will have a network connection ready and will be at the conference room at 10:00 A.M. to assist with the setup.

3/19 9:42 The 891 CSU technician was called to advise him it was time to go to 891/3098. He didn't answer so a voice mail was left so that he would be paged.

One person headed for 891/3098 and one went to 880/X10. Upon arriving at 891/3098, the 891 CSU technician was already there. He connected the test laptop to the network and configured the network parameters. Network access was tested and it worked correctly. Meanwhile, at 880/X10, the 880 CSU technician connected the other test laptop to the network and configured the network parameters. When the laptops were ready for testing, one laptop successfully accessed the hard drive of the other. The department server was also successfully attached to each laptop. Browsing through the network was successful. And the Sandia home Web page was successfully displayed in the web browser.

The network access part of this test went smoothly as expected. However, the process required for establishing this network access was time-consuming and sometimes difficult. A total of ten telephone calls were required to make the arrangements. Some guessing was necessary regarding whom to call for service.

A knowledgeable computer and networking user completed these trials under the guise of a "typical" user; one who knows little of netmasks, DNS servers, subnets, etc. Although there was assistance provided by several organizations, notably the CSU, this trial proved that a request for network access in a conference room could be quite painful. This service is not currently well supported from a customer viewpoint. There are simply too many hoops to jump through to obtain a service that could be provided in a user-friendlier manner. For example, the second trial required a total of ten telephone calls over the course of two days and two support people who provided on-site directions for connecting to the network.

Mobile computing can dramatically improve the usefulness of networking in a conference room. A user should be able to walk into a conference room, connect a computer to a network port, boot the computer, and begin communicating. It really can be that easy. The following section describes some methods that can make this a reality.

Methods for Providing Mobile Wired Networking

Three methods of providing mobile wired networking were studied for this project: DHCP, Mobile IP, and VLAN. Each method provides a means for allowing a mobile computer to connect to a network at different access points with no re-configuration requirements, no networking knowledge, and no assistance from network "gurus."

DHCP

DHCP, Dynamic Host Configuration Protocol, provides a mechanism for a computer to automatically be assigned a network address from a DHCP server. This is in contrast to the age-old method of manually assigning an address to a computer, keeping records of computer address assignments, and changing IP addresses when changes are necessary.

DHCP can alleviate the administrative burden of address management, but for the purposes of this project, the most desirable attribute is the ability to automatically assign a network address as a computer moves from one network to another. This gives a mobile user transparent mobility. A DHCP server takes care of the administrative functions automatically. The user is not required to make any changes to the mobile computer; rather the mobile computer automatically requests a new address when making a move and a DHCP server responds with an address assignment.

Providing a network address to a computer is a process analogous to renting an apartment. An apartment is typically leased to a tenant for a given period as determined by the landlord. If the tenant wishes to remain in the apartment beyond the lease period, the lease may be extended. If the tenant wishes to leave at the end of the lease, the lease is terminated and the apartment is made available for a new tenant. In the same manner, DHCP will lease an address to a computer for a period of time as determined by the administrator. If the computer needs the address beyond the lease time, the lease can be extended. If the computer no longer needs the lease, the lease is terminated and the address is made available for a new computer.

DHCP was created by the Dynamic Host Configuration Working Group of the Internet Engineering Task Force as an extension of the Bootstrap Protocol (BOOTP)[2]. It is an "Internet Draft Standard Protocol" in which "*The IESG is actively considering this protocol as a possible Standard Protocol. Substantial and widespread testing and comment are desired. Comments and test results should be submitted to the IESG. There is a possibility that changes will be made in a Draft Standard Protocol before it becomes a Standard Protocol.*" [3] Because DHCP is an Internet Draft Standard Protocol, it is widely available and on track to become a Standard Protocol. There are several commercial products available[4], and large organizations are beginning to take advantage of the technology.[5]

Mobile IP

Mobile IP[6] is a transparent method of allowing a network client to move from network to network while utilizing a fixed IP address. A client is assigned a fixed IP address as with conventional IP addressing. Therefore, the client can always be associated with this address. When the client moves to another network, a network agent must determine the identity of the mobile client and forward data from the client to the home agent who acts on behalf of the mobile client in the mobile client's home network.

The transparent nature of Mobile IP is the result of the fixed address of the client. A user has no need for reconfiguring a machine when moving to a new network. One simply moves to a new location, plugs in, and starts communicating. This is a very desirable trait of Mobile IP. However, a network agent must be present in each network segment that a Mobile IP customer wishes to utilize.

However, Mobile IP is an "Internet Proposed Standard Protocol. *These are protocol proposals that may be considered by the IESG for standardization in the future. Implementation and testing by several groups is desirable. Revision of the protocol specification is likely.*"[3] Because Mobile IP is a Proposed Standard Protocol, it is not yet widely available and no known commercial products exist. Although it seems to hold tremendous promise for allowing network clients to roam, the project team elected not to pursue this as a solution, electing to stay with commercial off-the-shelf solutions only.

Virtual LAN

Virtual LAN (VLAN) technology is a proprietary technique that assigns a network switch port to a network by way of software. This technique can be utilized for mobile network clients by dynamically changing the virtual LAN assignment for the port that a mobile client has plugged into. As with Mobile IP, a mobile client using VLAN technology has a fixed IP address and no configuration changes are necessary when moving.

A disadvantage with VLAN technology is its proprietary nature. Each manufacturer of network switching equipment has developed unique VLAN software. There is no inter-operability amongst different implementations. Although there are attempts to standardize VLAN technology, the current climate requires that a VLAN implementation be restricted to a single vendor. The project team elected not to pursue VLAN as a solution to mobile computing due to the proprietary nature. In addition to being restricted to a single vendor, any implementation would require switching gear that is VLAN capable. Older shared-media network gear such as Ethernet hubs would have to be replaced with newer equipment. Therefore, any areas that are not currently serviced by VLAN capable hardware would be unable to utilize mobile computing until the hardware was upgraded.

DHCP is Chosen for Mobile Wired Networking

The requirement of transparent mobile computing can be met by any of the aforementioned methods. However, the project team was looking for a solution that was both realizable and open. This would allow the solution to be put forth with minimal effort and not be locked into any one-vendor solution.

As Mobile IP is still in a research phase, this was not an acceptable choice. Although Mobile IP promises to provide transparent mobile computing using IP addressing, there are no known commercial products as of yet. This could change in the relatively near future as continued experimentation is conducted and the Mobile IP Internet Draft is finalized.

VLAN technology is currently proprietary, although there is work underway by Cisco Systems, Inc. and other manufacturers of networking equipment to standardize. The project team elected not to select a technology that placed heavy dependence on a single vendor. In addition, the technology is not available everywhere within the Sandia network, resulting in some areas being denied mobile networking opportunities until some time in the future.

DHCP is widely available as a product, is an open Internet standard, and can be applied to virtually any part of Sandia's network. For these reasons, the project team chose to implement mobile wired computing using DHCP.

Which DHCP to use?

Several DHCP products were evaluated and QIP from Quadritek Systems, Inc. was chosen for this project. See Appendix A for details of the evaluation and selection process. Due to time limitations, the evaluation process was informal. The chosen product is not necessarily the best or most appropriate product on the market, but was selected in order to proceed with the project. Other DHCP products could certainly be utilized in the future, if they were found better suited.

Security Issues Related to DHCP

There are at least three important security issues related to DHCP; what computer had a given IP address at any time, who made changes to DHCP configurations, and who can make changes to DHCP configurations. These issues are covered in detail in Appendix B for the QIP solution. The following section describes general solutions to each of these issues.

What computer had a given address at any time?

A particular IP address might be involved in some suspect activity such as accessing an inappropriate Web site. In a static IP world, it is simple enough to look up the computer associated with the given address as kept in some database such as DNS, then track down the computer and owner. However, in a dynamic IP world using DHCP, this is a more challenging process. Permanent records don't exist that show the address/computer relationship, since this relationship can and does change with DHCP.

Resolving computer/address relationships in a DHCP world can be done if a record is kept of both active leases (IP address assignments) and past lease activity.

Active Leases

A typical DHCP server product will display active leases. This would probably be the first place to check when attempting to discover a computer/address relationship. The display that shows active leases shows all leases that are currently assigned. This will show the MAC (Ethernet) address associated with the IP address. This is only part of the puzzle. Now we know a fixed attribute of the computer, since the MAC address is permanently associated with the computer[b]. We can then look up the MAC address in another database that keeps a MAC/owner relationship and find the owner of the offending computer.

This approach assumes that the offending computer is "DHCP active." That is, it has an active lease. However, what if the event occurred some time ago, say last month. The active lease list may or may not show the same IP address/MAC address relationship that existed one month ago. The lease time becomes a factor here. If the lease time is less than one month, then it is possible for the offending computer to have been assigned a different address one month ago. If there was a very low level of network activity from the computer, then the lease for the IP address may have expired. The original leased IP address may now be assigned to another computer on the network. With a long lease time, such as one month, the active lease information would probably reflect the same address as one month ago. However, if this is not the case, then another approach must be taken in order to discover the computer/address relationship. We must look at the past lease activity.

Past Lease Activity

Perhaps we determine the computer/address relationship from the list of active leases and find that the lease has only been active for a few days. We therefore can not assume that this relationship involves the alleged incident. In this case, we must search the DHCP server event log for DHCP activity, looking for the IP address in use at the time of the incident. We know the precise time the event took place. Therefore, the IP address active at the time of the incident should be identifiable in the log. For example, say an incident occurred on January 3rd involving address 134.253.4.180. We check the event log for DHCP lease activity. We find that no lease activity occurred on January 3rd for 134.253.4.180. However, on January 1st, a lease was renewed for 134.253.4.180 for MAC address 00-60-97-85-c8-03 for a one week period. Therefore, the address was leased at the time of the incident, we can now look up the user associated with this MAC address, and the culprit has been identified.

Mobile Address Identification Compared to Fixed Address Identification

The mobile address identification processes described above are rather convoluted compared to the straightforward method of identifying a fixed address computer. To identify a fixed address computer, one simply needs to look up the fixed address in a database such as Sandia's Network

[b] This is not true in all cases. If a network interface card is changed in a computer, the MAC address will change, and documentation should be updated. In addition, some workstations are capable of changing the MAC address through software.

Information System (NWIS). The address is directly associated with a user. Note that this does not identify the actual user involved with the computer; it only identifies the person associated with the computer.

In a mobile computing environment, where the address identification processes are more difficult than with fixed addresses, the additional administrative burden should be well offset by the simplified address management brought by letting a DHCP server manage addresses. Moreover, the task of identifying a user of an address is non-routine. The process is typically only utilized when inappropriate network activity is detected.

Who made changes to DHCP configurations

Keeping a record of changes to DHCP configurations is important in order to understand when and why changes were made. For example, we might be experiencing a problem with getting DHCP service to some nodes. In checking the change records, we find that an administrator made a change the previous day that resulted in an incorrect configuration that caused a failure. We can then roll back the change and attempt to determine how to make the proper configuration change based on what the administrator was attempting to do.

Who can make changes to DHCP configurations

Administrative control is important in a DHCP environment in order to maintain consistent configurations and prevent incorrect configurations that could cause significant chaos. A DHCP server must at minimum have a primary administrator account that can be protected by username and password. Additionally, secondary administrators may be helpful for dealing with tasks such as managing subnets and generating reports. These secondary administrator accounts must also be username and password protected with configuration rights that are a subset of the primary administrator.

Handling DHCP Requests Across the Network

In its basic implementation, DHCP will service clients on a local network segment only. DHCP requests from clients are broadcast only to the local network segment. This is due to routers not passing broadcast packets to other network segments. Any DHCP server located within the network segment should respond to the client request. However, any client located in a network segment that does not have a DHCP server will not receive a response from the DHCP server located in another network segment.

Rather than putting DHCP servers in every network segment that services DHCP clients, routers can be configured to pass DHCP requests on to DHCP servers as illustrated in Figure 4. This allows for a limited number of servers placed strategically within the network, as long as they and the network segments are of sufficient capacity to serve many network segments.

For Cisco Systems Inc. routers, a parameter called IP_HELPER_ADDRESS[c] is set for each router interface that needs to forward DHCP requests. This parameter will actually cause the router to pass all UDP broadcast packets to the specified destination. Included within the UDP broadcast packet types are DHCP packets. The broadcast packets are modified by the router to become unicast, which means they are sent to only one destination rather than all destinations. This has the disadvantage of increasing traffic on the network segment with the DHCP server.

The IP_HELPER_ADDRESS can be specified multiple times per router interface, so that more than one server can receive the DHCP requests. This is significant in order to have backup servers that can receive the requests if the primary server is unavailable.

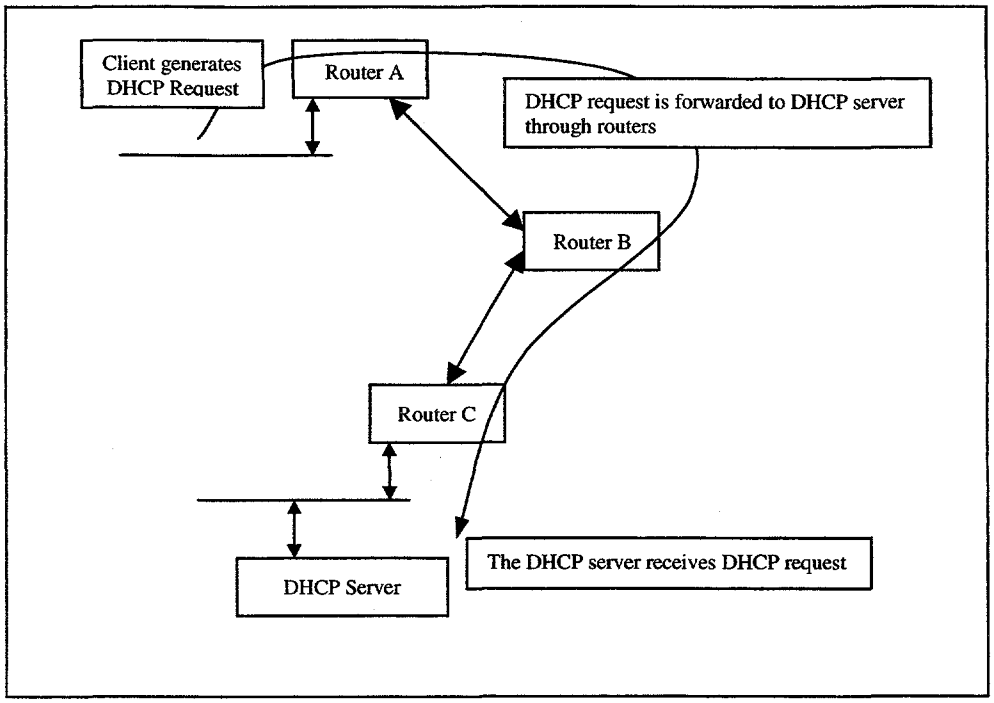
Client generates DHCP Request
Router A
DHCP request is forwarded to DHCP server through routers
Router B
Router C
DHCP Server
The DHCP server receives DHCP request

Figure 4. DHCP Packet Routing

[c] **ip helper-address (from the Cisco IOS 11.3 manual)**
To have the Cisco IOS software forward User Datagram Protocol (UDP) broadcasts, including BOOTP, received on an interface, use the **ip helper-address** interface configuration command. To disable the forwarding of broadcast packets to specific addresses, use the **no** form of this command.

Missing or Conflicting Network Services Confronting Mobile Computing

We have been living in a fixed-address, fixed-location world. We want to live in a mobile world; access information from anywhere, anytime. What are the network ramifications of moving a computer to a different location and expecting transparent access? Do any network services deny service to a roving computer? Are there procedures or processes that conflict? In a word, yes.

Name Service

The primary network service that is in conflict with computer mobility is DNS. DNS is involved with most every network connection, working silently in the background. Users are routinely unaware of the actions of DNS, but this powerful tool simplifies our lives by allowing us to use computer names rather than network addresses, with DNS doing the name-to-address translations. For example, the name www.sandia.gov is far easier to use than remembering 132.175.1.4. By entering www.sandia.gov into a Web browser, DNS will silently translate this into an IP address and the Web server is contacted. There is a hierarchical tree of DNS servers as shown in Figure 5. When a client contacts a DNS server, the DNS server uses its local database of resolved names and IP addresses to attempt to resolve the supplied name. If the name cannot be resolved from the entries in the local database, the request is forwarded to the next higher server in the hierarchical tree of DNS servers. As shown in Figure 5, the Sandia name server goes to a root name server sitting at the top of the Internet DNS name server tree. This name server knows of all other name servers and returns the address of a name server that should have the requested information.

DNS has traditionally been a static database of computer names and addresses. A computer's name and address are entered into DNS, typically via a manual process, and they stay in DNS until manually removed. Of course, some computers never get into DNS and remain "unlisted." A typical home computer accessing the Internet has no need for a DNS entry since no other computers need to access a service on this computer. However, servers certainly need to be in DNS so users can access services such as Web sites by name, rather than by address. This has been a sufficient arrangement because most computers that need a DNS entry have been static in nature; neither the name nor address changes frequently.

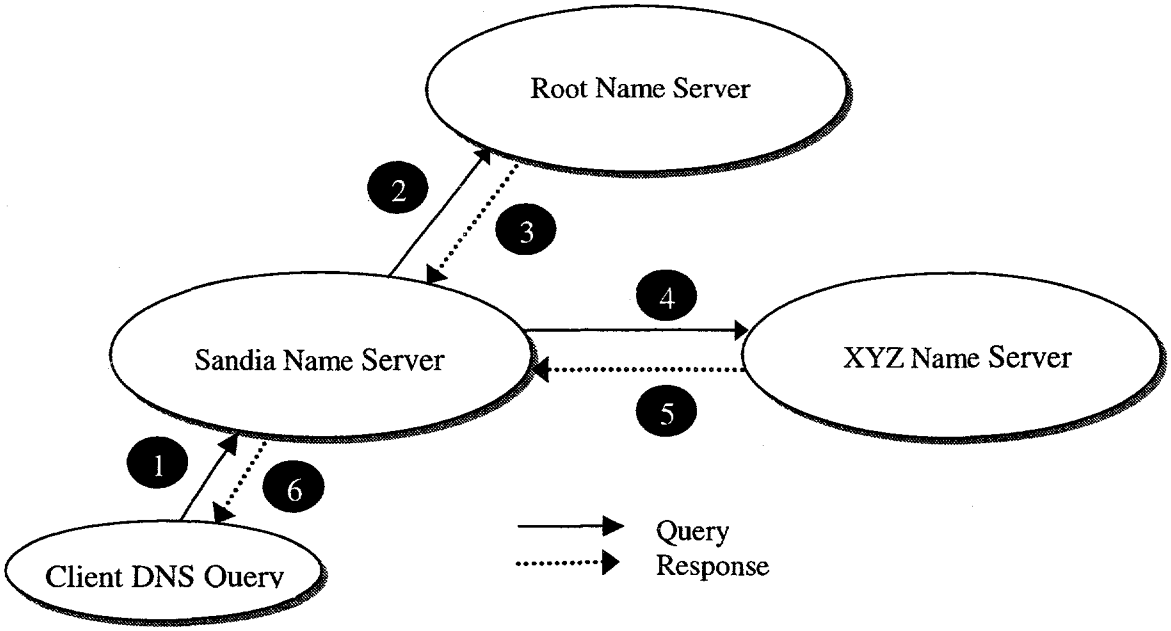

1 A client computer requests a DNS name-to-address translation for host xyz.company.com

2 The Sandia name server does not have the address for xyz.company.com. The data has not been cached since no previous look-up has been attempted. The query is therefor passed to a root name server.

3 The root name server knows the name server for company.com and returns the address for the XYZ name server to the Sandia name server.

4 The Sandia name server queries the XYZ name server for the address for xyz.company.com.

5 The XYZ name server sends the data back to the Sandia name server.

6 The Sandia name server then sends the data on to the client, completing the query. The client now can contact xyz.company.com by address. The data is also cached in the Sandia name server to facilitate future look-ups.

Figure 5. DNS Name Resolution Process

The conflict arises out of the desire to provide mobility to users by using dynamic addressing, while using a naming service that is designed for a static world. While there are ways to provide mobility while maintaining static addresses, DHCP changes the address of a computer when moving to a different network. A DNS server, filled with addresses that are not expected to change dynamically, will not keep up with a computer that changes addresses "on the fly." This can be mitigated by using a dynamic name server, called DDNS (dynamic domain name server). With DDNS, when a computer receives a new address, DDNS is updated immediately with the new computer name/address pair. Therefore, any other service or user that needs to contact or refer to this mobile computer will receive the current name and address data.

Note that there is another name service called WINS (Windows Internet Name Service) that also provides name to address resolution. This is a Microsoft product designed to provide computer addressing and naming without reliance on DNS. However, this is a dynamic naming service and therefore does not conflict with mobile computing.

No DNS Entry = No Service (Sometimes)

Some applications require a client computer to be registered in DNS. For example, the Sandia Usenet news server will not respond to a computer that is not in DNS and therefore service is denied. Some Internet FTP servers will not respond to anonymous file transfer requests unless the requesting computer is in DNS.

This poses a problem for a mobile computer that is using DHCP. The dynamic addressing nature of DHCP clashes with the static addressing nature of DNS!

DHCP and DNS – How Can They Work Together?

DHCP assigns IP addresses to network clients without manual intervention, thus reducing the personnel required when new clients are established or moved throughout the network. However, this is in direct contradiction to our current addressing strategy where an IP address is statically assigned to each computer and entered into the Domain Name System (DNS) by an administrator. The fallout from having statically assigned IP addresses is that the DNS is also a static function. Once a computer is assigned an IP address, that address is placed in the DNS and remains there until removed by an administrator.

DNS provides a mapping between IP addresses and host names. We can look up a host by name rather than having to remember a string of digits comprising an IP address. Conversely, given an address, we can look up a host name. This is mostly transparent to users. The magic happens in the background, unless one is doing troubleshooting and needs to know an address or host name given one or the other. To go to a Web site such as www.sandia.gov, we simply enter the name. Our computer then queries DNS for the address associated with this name and makes a connection to this address. In some cases, this works in reverse. When attempting to connect to some applications, the application will look up the host name given the address, to see if this computer has been identified in DNS, thereby making it a "legitimate" node on the network.

In utilizing DHCP, IP addresses are assigned for a specified (but renewable) time interval and are then returned to the pool of available IP addresses when no longer needed. Clients may or may not be assigned the same address by DHCP. In cases where computers are not moved and are connected regularly to the network, assigned addresses will most likely not change over the course of days, weeks, or even months. However, when a computer is moved to another location, most likely the address will change due to connection to a different network segment. How do we manage the existing DNS with computer addresses changing? This can not be a manual process. If a user accesses the network with a laptop at the user's desk by utilizing DHCP, then the user takes the laptop down the hall to a conference room and accesses the network again through DHCP, there should not be any requirements for administrative changes to update the network information.

A solution to this problem is to establish a dynamic DNS (DDNS) that services DHCP clients. We can then link the static corporate DNS servers to this dynamic DNS server in order to make the DHCP clients known to the rest of the network.

With DDNS, a client computer that is assigned an IP address through the DHCP server has the IP address information automatically loaded into the DDNS server. Any computer on any corporate network as well as on a network anywhere in the world can then resolve the name or address of this DHCP client.

How is this done?

1. When a DNS query takes place on the Internet, looking for a computer in Sandia's IRN, the Internet DNS points to a DNS server in the IRN.

2. The IRN DNS server then looks up the address and returns the result to the Internet DNS server.

3. If the lookup is for an IRN DHCP client computer, the query is passed to the IRN DNS server, which in turn passes the lookup to the DHCP DDNS server, which passes the result back to the IRN DNS server, which passes it back to the Internet DNS server.

In the long run, the corporate DNS servers could be replaced with a set of dynamic corporate servers, eliminating the need for a special DDNS server. This would eliminate the need for additional DNS servers that exclusively handle dynamic DNS.

Potential Name Conflicts

A dynamic naming service (DDNS) opens a door to potential conflicts with names with respect to the static name service (DNS). What is to stop a user from naming their computer with the same name as one that already exists in DNS, thus having the same name in both DDNS and DNS? With the two systems running independently but cooperatively (i.e.: when DNS can't resolve a name, it can try DDNS), this could spell disaster.

But this problem is completely avoided by having the DDNS servicing a sub-domain of the primary Sandia domain sandia.gov. For example, we might set up a sub-domain called dhcp.sandia.gov. Any computer that is dynamically addressed in this sub-domain will be guaranteed a unique node name due to the inclusion of the sub-domain name. If there exists a computer called joe.sandia.gov in the primary domain, and a DHCP client computer is configured as joe, it will be named joe.dhcp.sandia.gov in DDNS. Therefore, there is no conflict with joe.sandia.gov.

However, that is not to say there won't be some confusion. Asking for the address of joe could return joe.sandia.gov or joe.dhcp.sandia.gov, depending on the sub-domain location of the inquirer. If node bill, located in sandia.gov, asks for the address of joe, bill will receive the address for joe.sandia.gov. If node bill, located in dhcp.sandia.gov, asks for the address of joe, bill will receive the address for joe.dhcp.sandia.gov, if this node exists. Otherwise, bill will receive a message saying joe does not exist. The workaround for this confusion is to use full names such as joe.dhcp.sandia.gov. This is mostly an issue when users look up computer names and addresses in DNS. When computers look up names and addresses, they typically use full names.

How does DDNS handle Its Own Potential Name Conflicts?

In addition to potential name conflicts between DNS and DDNS, how does DDNS handle potential name conflicts within itself? Due to its dynamic nature, DDNS can see computers with duplicate names, or the same computer moving about the network. The end result is having the same name associated with different IP addresses.

A Computer Moving About the Network

When a computer moves to another network segment, and through DHCP, receives a new IP address, DDNS is updated with the new name/address pair. Any existing name/address relationship for this computer is purged. Therefore, there can be no ambiguity regarding the name/address relationship when a computer moves and receives a new address.

Duplicate Computer Names

When two computers are configured with the same name, DDNS takes a "first come, first served" approach. Once the first computer is registered in DDNS with the user-assigned name, any subsequent computers with the same name will be registered in DDNS with a default name that has been preassigned to the address by the DHCP server and is guaranteed to be unique within the DDNS name space. This guarantees that the computer can utilize a service that requires DNS name registration. However, the user is unaware that the chosen computer name was rejected and replaced by some unique default string generated by the DHCP server.

NWIS

The Network Information System (NWIS) is a corporate database containing network details for computers connected to Sandia networks. As with DNS, NWIS is based on the notion of fixed network addresses. A given computer has one or more addresses (in most cases only one) and each address is associated with the corresponding network. For example, computer

SAGW040.SANDIA.GOV is associated with address 134.253.4.12 in NWIS and is therefore a member of network segment 134.253.4.

NWIS is another service that conflicts with mobile computing. If machine SAGW040.SANDIA.GOV were to participate in mobile computing using DHCP, NWIS would not be able to represent this function, since it would want to associate a particular address with this computer. Each computer in NWIS that is connected to a network is documented as a LAN connection. Mobile computers do not have a particular LAN connection; rather they have any LAN connection.

To meet the requirement that all machines be documented in NWIS, a mobile machine could be entered with just machine data and make no association with a network. However, one important piece of information would be missing; the MAC address, which is stored in the LAN connection. This address, also known as the Ethernet address or hardware address, is an important element for determining the link between a particular machine and an IP address assigned by DHCP. See Appendix B for details on how a MAC address can be used to link a DHCP client to an IP address. Mobile computers could be completely identified in NWIS by changing NWIS to do the following:

- When establishing a LAN connection for a machine, associate the machine with DHCP or some other designator, indicating this is a DHCP client and therefore has no fixed IP address or network association.
- All other LAN connection information can be maintained as is, including the important MAC address.

With this change to NWIS, all mobile computers can be grouped together for reporting and for database queries. For example, all mobile computers can be reported; or all those belonging to a particular department; or all those belonging to a particular person. Moreover, with the MAC address included, machine owners can be identified when network activity warrants an investigation.

DHCP Trial

Mobile wired networking, and in particular DHCP, was put to the test in an office environment. Several office computers were configured to be DHCP clients. Over a period of approximately six months, these DHCP clients were utilized on a daily basis.

DHCP clients for this test included typical desktop systems, laptops, and an H/PC. The desktop PCs were all running Windows NT 4.0. Some laptops were running Windows NT 4.0 and some were running Windows 95. The H/PC runs Windows CE 2.0. In all cases, the clients successfully obtained and maintained leases.

A DHCP server was established on a Dell Precision 410 PC running Windows NT Server. The Precision 410 is a 400Mhz PC configured with 64MB of RAM and a 4Gbyte hard drive. Quadritek's QIP was used as the DHCP software.

A backup DHCP server was put in place to test redundancy capability. Quadritek's DHCP redundancy scheme is to have the backup server's DHCP function disabled, but the server still talks to the primary server on a continuous basis to determine the state of the primary. If the primary does not answer within some configurable time interval, the secondary assumes the primary is down and enables its DHCP server. The secondary will continue to attempt to communicate with the primary and will disable itself when the primary returns to service. The primary and secondary servers exchange address information periodically in order to remain synchronized.

The redundancy feature was tested and proved to operate as expected. The primary server was taken off the network to simulate a failure. The secondary server, being unable to communicate with the primary, then took over the DHCP function. This was tested by renewing a lease on a computer that had an active lease assigned by the primary. The lease was renewed with the same address and other parameters as were originally set. Then the primary server was reconnected to the network. After waiting a few minutes for the primary and secondary servers to synchronize, the primary server's active leases were viewed. The primary correctly showed the renewed lease that had been given by the secondary server.

In addition to the DHCP server, Quadritek's DDNS name server was also tested. A sub-domain called lip.sandia.gov was created. All DHCP clients in the test were successfully entered into the name server and the lip.sandia.gov domain as they received a lease. The DDNS server also operated successfully in translating names to addresses for outbound connection requests coming from the DHCP clients. However, the server was not configured to provide name to address resolution for inbound connection requests. This would require establishing a pointer in Sandia's primary name server that points to the test name server and it was felt that the production system shouldn't be touched in order to prevent any service disruptions. Had the pointer been established, it would enable computers outside Sandia to resolve the addresses of the DHCP clients in the test bed, just as can be done with Sandia's production name servers.

Over the testing period, no significant problems were encountered. The primary server ran solidly for the entire test for both DHCP and DDNS services. The secondary server, other than during redundancy tests, never needed to take over for the primary. All test clients behaved well.

However, a significant lesson learned from this trial was that occasionally the DHCP server would receive DHCP-like requests that it should have ignored but didn't. This caused the originator of the requests to receive unexpected responses. Specifically, a node was broadcasting BOOTP messages (a protocol that predates DHCP and is often included in DHCP servers) in order to boot diskless from another machine. The DHCP server was intercepting these messages and attempting to respond, which confused the machine trying to connect to its boot partner. The workaround for this problem was to establish a pool of client MAC addresses in the DHCP server that were eligible for DHCP service. Any machine in this MAC pool could be serviced by the DHCP server. Therefore, the machine attempting to boot diskless would not be bothered by the DHCP server since the machine's MAC address was not included in the MAC address pool. This may very well be a good technique to use in any future implementation of DHCP. A MAC address pool prevents unwanted or unauthorized users from receiving DHCP service, and keeps the DHCP server from conflicting with other services. On the other hand, this adds to the administrative tasks for maintaining DHCP, but the MAC pool update process could be automated to mitigate this additional burden.

Wireless Technology

Wireless technology was studied with the hope of providing an even greater means of mobile computing: no wires to tie us down. Unfortunately, there are a number of security issues to be resolved before wireless technology could be deployed at Sandia. The results of the study are documented below; however, a decision was made not to proceed beyond an initial study phase of wireless technology until the security issues were addressed properly.

Wireless networking technology is experiencing considerable growth. Products are now available that are both affordable and practical. While there are several issues of usability that remain, such as security of data transmissions, bandwidth, and range, wireless is certainly a technology to be reckoned with.

Security issues involving wireless networking include the need to encrypt data since the data is broadcast through the air and the need to authenticate to ensure users are legitimate. These requirements are similar to extending the IRN from one area to another (encryption) and remote access to the IRN (authentication).

Beyond the problem of securing wireless data transmissions, the most glaring drawback to wireless is bandwidth. Typical bandwidth is in the range of 1 to 2Mbps, roughly 10 to 20% of a 10Mpbs Ethernet, while some new technology is reaching as high as 10Mbps. Wireless bandwidth will probably always lag that of wired bandwidths. For example, Ethernet is now routinely running at 100Mbps. It will be some time before wireless can achieve anything near this rate. Also, as wireless data rates increase, effective range tends to decrease. A 10 Mbps wireless network will have a shorter range than a 2 Mbps wireless network.

There are two primary methods of transmission of data; Frequency-Hopping Spread-Spectrum (FHSS) and Direct-Sequence Spread-Spectrum (DSSS). FHSS, as the name implies, changes frequency in a pattern known to both the transmitter and receiver. Data is transmitted in a short burst on one frequency, then in a short burst on another frequency, and so on. DSSS spreads the transmitted signal over a wide frequency range, using a spreading code known only to the transmitter and receiver. The only significance of being aware of these technologies is that some manufacturers use FHSS and some use DSSS. There are advantages and disadvantages to each technology. At this time, FHSS is more popular, but DSSS is beginning to challenge FHSS as wireless transmission rates increase. DSSS is more suited to higher bandwidths but more difficult to implement.

Wireless data transmission does not have licensing requirements. The typical frequency is 2.4Ghz.

A number of corporations and universities are currently using wireless technology, and in the San Francisco Bay Area, there is an ISP that is providing connectivity via wireless technology.

Information on wireless technology was gathered from Web searches, by attending an Albuquerque telecommunications trade show, the Networld+Interop trade show, and by contacting vendors directly. Some of the vendors identified are noted below.

Wireless LAN products have historically been designed with proprietary technologies, making interoperability impossible among vendor product lines. However, for the past several years, the IEEE has been working toward a wireless standard called 802.11, which was approved in June 1997. This will improve the chances that one vendor's product will work with another's. However, issues remain that are not addressed by 802.11. One issue is roaming. This is still a proprietary feature that is not guaranteed to work when roaming between cells supported by different vendors.

Some Vendors

The Breezecom wireless communications system [www.breezecom.com] includes the BreezeNet Pro family of equipment which includes wireless hubs (around $695 for a single port and $1195 for 4 port station) and PCMCIA adapters (around $565) (see Figure 6). The system is RF based and uses FHSS in the 2.4GHz to 2.4835GHz band. The company claims 1000m range in an unobstructed environment and 60-200m in office environments. Security of the communications is based on the use of FHSS (eavesdropping is made difficult by the constant frequency shifting). Data rates are in the 1-3 Mbps range.

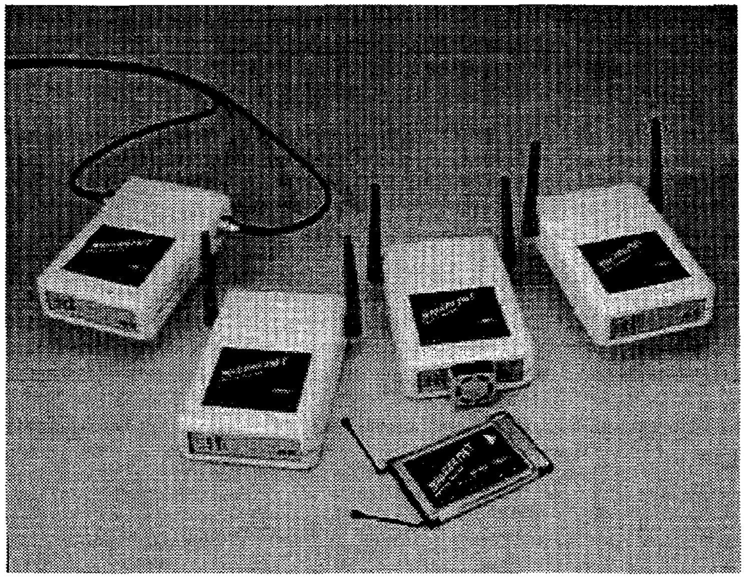

Figure 6. Typical Wireless Equipment from Breezecom

The Proxim [www.proxim.com] system also operates in the 2.4 GHz range using FHSS and has similar data rates. Both a PCMCIA card (with two antenna styles) and an ISA bus connection are available. Security of the communication is based on the use of FHSS transmission and other features. The system is used by the Pacific Stock Exchange for communicating the highly confidential information about trading transactions. Proxim believes that their equipment is safe around medical devices, and the equipment is used by a number of hospitals - even in ICU's.

The wireless equipment from Lucent [www.wavelan.com] uses Direct Sequence Spread Spectrum (DSSS) technology. Lucent also has an optional encryption chip for added security.

Several vendors use infrared (IR) technology for a wireless communication. The IR technology is limited to line-of-sight range between the transmitter and the receiver. The IR technology might work well where maintaining line-of-sight is not a problem such as in a conference room setting or in an office area with low dividers; however, the RF technology would be much more flexible when used in a roaming mode.

Wireless Applications

A campus-wide wireless mobile Internet network communications architecture is being developed by Purdue University. The prototype system is referred to as the Crosspoint design project.[7]

Wireless technology has the potential for providing a number of capabilities to the roaming worker. These include roaming within a building or a larger outside area using wireless technology such as IR or RF. An in depth study of these technologies was deferred until the security issues of encryption and authentication could be studied and resolved. A future project may attempt to explore the technology for wireless LAN's (e.g. RF and IR) and to survey some industrial and government installations to investigate successes, failures, handling of security issues, etc. for a worker who is roaming around a building or an area while maintaining a connection to the local computer network.

Besides roaming capability, wireless networking can be applied to a stationary situation, such as a small building that has no network connections in place. A LAN can quickly be established with wireless stations and hubs to interconnect all computers in the building. An inter-building wireless link can then interconnect the building to other networks.

Simplification of the Remote Access Configuration Process

In exploring issues involving dial-up remote access, the project team discovered that the process of configuring a PC for dial-up access is difficult. There are a number of steps that must be followed precisely. The slightest misstep can result in a dial-up process that does not function, or does not function as desired. Sandia has long had a set of instructions available for making the dial-up configuration. These instructions were followed by several team members. While the instructions were found correct, the number of steps required to configure a machine for dial-up is daunting as a result of the complexity of dial-up networking. For example, a typical dial-up configuration requires 33 steps to complete.

This complex configuration process can be simplified by using a shareware program called Dialup Constructor. This program is designed to automatically configure a PC for dial-up access. It works with both Windows 95 and Windows NT. An administrator creates a dial-up profile, consisting of all parameters associated with a particular dial-up. This profile is installed on a user's PC by running a small program that automatically installs the dial-up profile. A user needs only to run this single program. No questions are asked and no decisions are required. At the completion of the installation, which takes only a few seconds, the PC is ready to dial (the configured dial-up is located in the Dial-Up Networking folder). If more than one dial-up is desired, such as for the IRN and the EON, the installation program can be executed for each profile.

Figure 7 shows the first of two administrative screens. On this screen, and the following screen shown in Figure 8, the administrator sets the parameters required for a particular dial-up configuration. This is done just once by the administrator to create each desired profile. Note that scripting is supported as shown at the bottom of Figure 7. Scripting is a dial-up technique that simplifies the dial-up process by automatically interacting with prompts for executing dial-up configuration commands. Profiles can be created for both scripted and non-scripted versions.

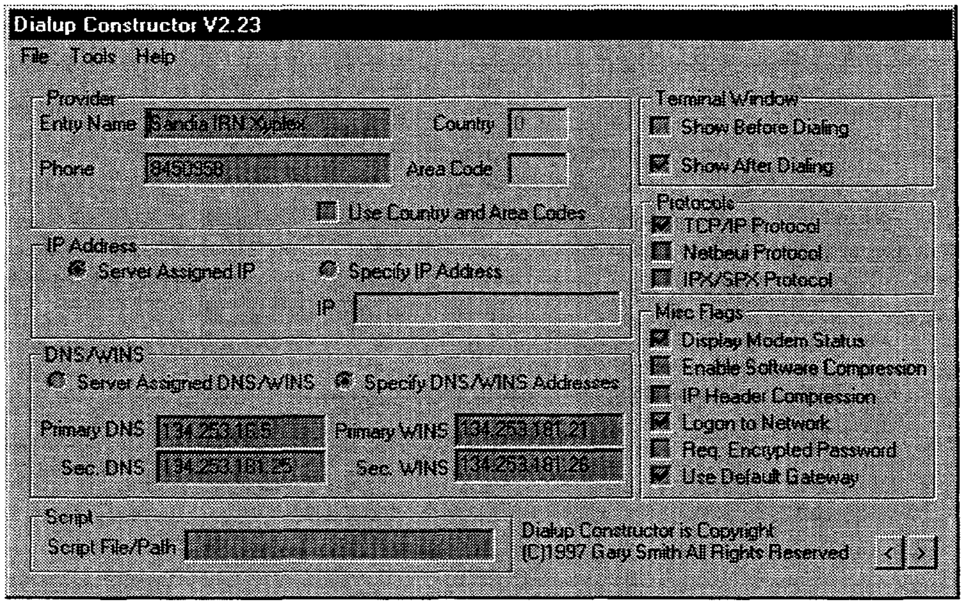

Provider
Entry Name: Sandia IRN Xyplex
Country: 0
Phone: 8450958
Area Code:
□ Use Country and Area Codes

IP Address
● Server Assigned IP ○ Specify IP Address
IP:

DNS/WINS
○ Server Assigned DNS/WINS ● Specify DNS/WINS Addresses
Primary DNS: 134.253.16.5
Sec. DNS: 134.253.181.25
Primary WINS: 134.253.181.21
Sec. WINS: 134.253.181.26

Script
Script File/Path:

Terminal Window
□ Show Before Dialing
☑ Show After Dialing

Protocols
☑ TCP/IP Protocol
□ Netbeui Protocol
□ IPX/SPX Protocol

Misc Flags
☑ Display Modem Status
□ Enable Software Compression
□ IP Header Compression
☑ Logon to Network
□ Req. Encrypted Password
☑ Use Default Gateway

Dialup Constructor is Copyright
(C)1997 Gary Smith All Rights Reserved

Figure 7. Dialup Constructor, First Screen

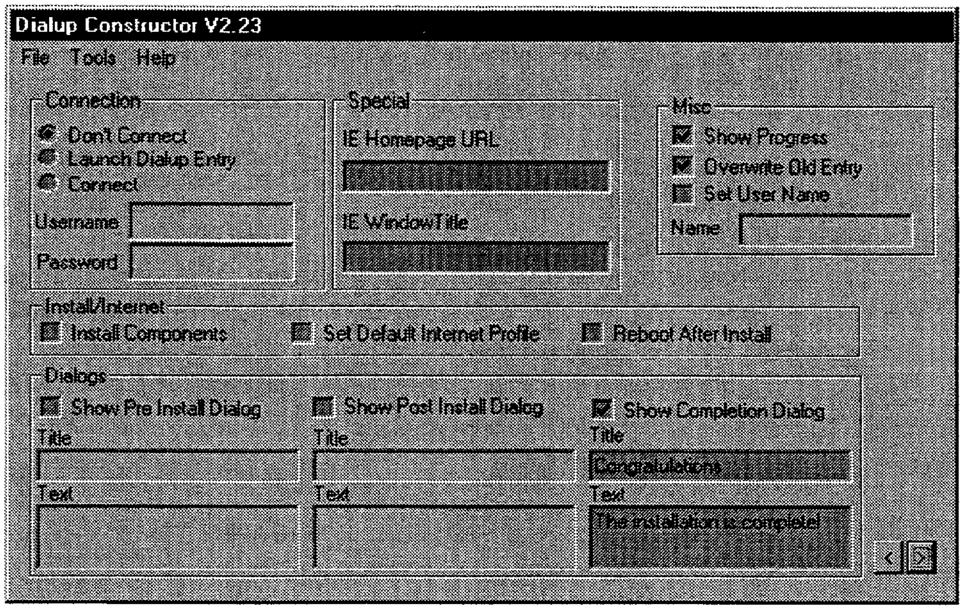

Connection
● Don't Connect
○ Launch Dialup Entry
○ Connect
Username
Password

Special
IE Homepage URL
IE WindowTitle

Misc
☑ Show Progress
☑ Overwrite Old Entry
□ Set User Name
Name

Install/Internet
□ Install Components □ Set Default Internet Profile □ Reboot After Install

Dialogs
□ Show Pre Install Dialog □ Show Post Install Dialog ☑ Show Completion Dialog
Title: Congratulations
Text: The installation is complete!

Figure 8. Dialup Constructor, Second Screen

This simplified procedure could be implemented in several ways. One method would be to give a customer a floppy disk that contains the requested dial-up profiles. By simply executing a file on the floppy, the user's PC is configured. Another method would be to make the files available on Sandia's intranet. A user could download the desired profiles and executable on a desktop PC and make a floppy for the laptop (if that is the desired destination). If the laptop is connected to the network, the files could be loaded directly onto the laptop and executed. And yet another method would be to distribute the files via SMS. If a user wishes to have dial-up networking installed, it will be done automatically. If not, the installation can simply be rejected.

An additional simplification process would be to utilize an installation program that could present to the user a set of profiles to select. The user could simply select the desired profiles and each would automatically be installed.

Three Scenarios for Location Independent Access

Three scenarios for Location Independent access were considered; mobile wired, mobile wireless, and virtual desktop. The mobile wired situation involves a computer that can be easily moved from one location to another, such as a laptop or hand-held, and connected by wire to a network. This is intended to be confined to the campus environment. A user can unplug from one location, move to another location, plug in, and have the same (in most cases) network access. The mobile wireless situation is the same as the mobile wired except that the computer is not wired to a network, but rather has a radio link to a network. This can and should be confined to the campus area, rather than extending radio links to other areas, and has the additional burden of security requirements due to the broadcast nature of the technology. The virtual desktop is a concept where a user can log on to any computer, see the same "desktop" or screen, and work as if at the office computer.

As discussed earlier, the mobile wireless and virtual desktop scenarios were deferred to later study due to security and availability constraints. However, having gained some knowledge in these areas, they will be discussed briefly.

Mobile Wired

The first scenario is for the mobile wired laptop, and this scenario has two cases, conference room access and office access. There is actually very little difference between the two cases as both require mobile-enabled computers, but conference rooms are challenging with respect to enabling a group to gather with laptops and conveniently connect to the network.

Suppose a project leader wants to hold a meeting with project members in a conference room to discuss some project data that is stored on a server. He brings his laptop to the meeting, plugs into a network port, powers up, accesses the desired data, and displays it to the group. How can this be achieved?

We need to provide a method of transparent network access from any location. Dynamic network addressing was the chosen method. This is accomplished using DHCP, Dynamic Host Configuration Protocol. This protocol will assign an IP address to a computer (called host or client) and therefore allow it to communicate on the network. When the computer is moved to another location, most likely resulting in connection to a different network, DHCP will assign a new IP address to the computer and communication can commence. This is totally transparent to the computer user.

To provide transparent network access to a conference room, a DHCP server must be established somewhere within the enterprise network. It does not have to be located in the same network that services the conference room. Therefore, a single server can potentially serve many conference rooms. For redundancy, a primary server and a backup server should be established, with the backup server located in a different network from the primary to avoid loss of both servers if the primary server network loses connectivity.

Router interfaces serving networks that have DHCP clients must be configured to forward DHCP packets to the DHCP server. As DHCP is employed, these router interfaces can be configured, rather than having to modify the entire router network.

A set of addresses must be established in the DHCP server that is available for leasing. This set would be comprised of groups of addresses for each network segment that is to provide dynamic addressing.

A DHCP client must be configured. This is a simple, brief process. For Windows NT, bring up the Network Properties dialog box by right clicking on the Network Neighborhood desktop icon, or selecting Network in the Control Panel. Go to the Protocols tab, select the TCP/IP Protocol, and either double click or select the Properties button. In the Microsoft TCP/IP Properties dialog box, as shown in Figure 9, select the IP Address tab. Select the button for "Obtain an IP address from a DHCP server". Click the OK button and reboot as prompted. This is the only change necessary to enable a computer to utilize DHCP. However, there are many other issues involving the client configuration for Windows 95 and Windows NT that are discussed in detail in Appendix D, which describes DHCP client configurations.

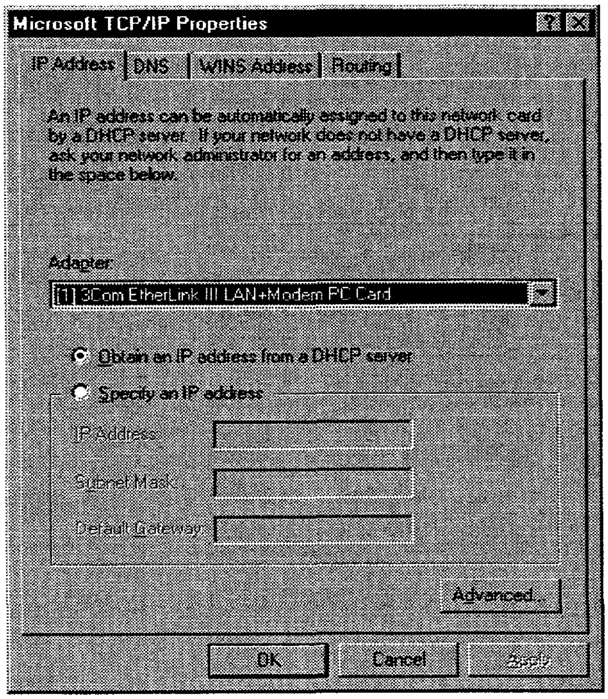

An IP address can be automatically assigned to this network card by a DHCP server. If your network does not have a DHCP server, ask your network administrator for an address, and then type it in the space below.

Adapter:
[1] 3Com EtherLink III LAN+Modem PC Card

● Obtain an IP address from a DHCP server
○ Specify an IP address

IP Address:
Subnet Mask:
Default Gateway:

Figure 9. Windows NT Client Network Configuration for DHCP

Conference Room Implementation

Conference room network access suffers from deficiencies:

- Network ports are not always wired to the network
- Network ports that are wired may not be enabled
- There may be no network ports or insufficient numbers or ports
- Network ports may be in inconvenient locations within the room
- Users may need expert assistance to set up networking
- User must make several phone calls to make arrangements for network access

These deficiencies can be overcome with mobile computing enabled with DHCP, and with some simple changes to conference rooms. Some conference rooms are not arranged in such a way as to facilitate a group that might wish to bring in several laptops and connect to the network. For example, in a large room with network ports along the wall, users would either have to gather along the walls, or run long wires from their seats. This scenario could be improved by having "network dispensers" strategically placed at each table. A table that seats three could have a dispenser containing three network ports and three power receptacles located in the middle of each table. On the other hand, receptacles could be located unobtrusively under the table at each seat. In both cases, users would only need a short network cable to plug in to the network, and a typical power cord would reach the power receptacle with ease.

For the purpose of the pilot study, a conference room was temporarily configured for mobile access. To avoid disturbing existing wiring, an unused fiber pair was utilized to establish a path to the DHCP server. A small, 4-port network hub was used for providing multiple ports. Based on an idea suggested in the User Input section, a device was constructed that houses four data ports and four power plugs, making it easy for a group of four to gather around a central point, such as a conference room table, plug in their laptops (data and power), and exchange data. This "network dispenser" is nothing more than a set of network and power receptacles housed in a box that can be conveniently located in a room.

The network hub used for the pilot demonstration is a "quick and dirty" solution only and is not recommended here for expanding port availability. The recommended procedure is to have a room wired with additional network ports. Network troubleshooting is simplified since there is no additional, potentially unknown, equipment strung on to the circuit.

With a DHCP server established and clients configured to utilize DHCP, mobile wired computing is enabled. However, to further enhance the performance of the system, a dynamic DNS server can also be established. This will provide dynamic name-to-address resolution for client machines as they move about the network.

An additional configuration for conference rooms would be to utilize VLAN technology in concert with DHCP to provide a single network segment that serves only conference rooms. This would provide the ability to offer short-term leases to conferences room where network access is typically short lived.

Office Implementation

Implementing mobile computing in offices is very similar to conference rooms, but physical access to network ports is normally much easier as most offices are already wired for each occupant. In addition, only one or two network connections are typically needed per office. Therefore, offices are already well suited to take advantage of mobile computing. The DHCP infrastructure must first be established by installing DHCP servers and configuring routers to forward DHCP packets. A machine configured for DHCP would simply be plugged into a network port in the office. An advantage to this office arrangement is that a user who wishes to do mobile computing within the campus area can easily move between the office and other locations such as conference rooms without making any changes to the computer configuration; DHCP all the way!

Switching between Fixed Addressing and Dynamic Addressing

There will surely be situations where a mobile computer user must operate with a fixed IP address. This might be the case where no DHCP service is yet available for the user's network segment, but the user would like to utilize a DHCP-enabled conference room. The switch from fixed address to dynamic address and back again is not a complex manual process, but can be daunting to the network-challenged.

Rather than having to know anything about address switching, an application called Netswitcher will do all the work. This program will allow a user to simply double-click an icon to change from one addressing mode to another. This operation is quite simple, but the configuration does require some networking knowledge. Each machine that uses Netswitcher must have its networking parameters set for the two addressing situations; one that is set for DHCP, and the other set for the fixed IP address and computer name for this computer.

Steps to Implement Mobile Wired Computing

For either a conference room implementation or a large-scale implementation covering both conference rooms and offices, the following steps illustrate the basic requirements for establishing mobile wired computing.

1. Establish Quadritek's QIP Enterprise Server in the network backbone. QIP Enterprise is comprised of a Sybase or Oracle database that holds configuration data and optionally a DHCP server and DDNS server. For an initial implementation, the Enterprise server could easily handle the DHCP and DDNS services. However, as demand increases, these services can be distributed to other servers. Place the server in a protected environment such as the computer annex in Building 880. Here the server would be maintained 24 hours per day. QIP runs on any Windows NT Server running on an Intel-based PC, or any UNIX server running Solaris 2.x, HP-UX 10.0x or AIX 4.1.2.

2. Establish a QIP DHCP server on the same machine running the QIP Enterprise database. As noted above, a separate machine can be established for running DHCP, but this in not initially necessary.
3. Establish a backup Quadritek QIP DHCP server somewhere in the backbone, but in a different network segment from the Enterprise server. This will provide protection from loss of the primary DHCP server either from the primary server going down or due to loss of the primary server network.
4. Establish a QIP DDNS server on the same machine running the QIP Enterprise database. As noted above, a separate machine can be established for running DDNS, but this in not initially necessary. The DDNS server can be configured as a subdomain of sandia.gov, such as lip.sandia.gov or dhcp.sandia.gov. Sandia's primary name servers, configured with a pointer to the DDNS server, could then provide name-to-address resolution of Sandia's DHCP clients from anywhere in the world. The DDNS server could be used in place of the existing static DNS servers and provide name-to-address resolution for all of Sandia's computers, both fixed and dynamically addressed.
5. If desired, a backup DDNS server can be established on the DHCP backup server.
6. Configure any router interfaces that are going to service DHCP client network segments by setting the IP_HELPER_ADDRESS parameter. This function will forward DHCP packets from the client to both the primary and secondary DHCP servers.
7. Configure user machines that are to utilize DHCP.
8. Mobile wired computing is ready.

There remains, in this implementation plan, the issue of whether or not to control access to DHCP by MAC address. Should any machine have access to DHCP service? Or, should a machine be required to register in some fashion before the service can be accessed? DHCP testing for this project was done using both methods. Each proved to be workable on a small scale. However, on a large scale, restricting access to DHPC service by MAC address creates an administrative burden by requiring a process for maintaining the MAC address list in the DHCP server. However, by not restricting access to DHCP service based on MAC address, the door is open for any machine to receive DHCP service, whether authorized or not. In addition, there is no incentive for a user to document their MAC address or keep the current documentation up to date since they could access the network whether or not their machine information is documented. This has a negative effect of making it difficult to track down users of IP addresses. Also, with the need to register a machine prior to gaining network access, NWIS would benefit from having all user machines documented.

When restricting access to DHCP service by MAC address, it is necessary to store in the DHCP server those MAC addresses that are to receive service. The DHCP server will then only respond to those machines whose MAC address is known. This could be thought of as "registering" one's MAC address prior to obtaining service. If this method is chosen, there are a number of techniques for getting the MAC address registered. NWIS would certainly be the primary repository of MAC addresses, just as it is currently. However, as pointed out earlier in this document, the current version of NWIS can not store MAC addresses without an IP address, and this would have to be modified in order to link a MAC address with a particular machine and not a particular network. An automated process could be developed that takes all the MAC addresses stored in NWIS and routinely updates the DHCP server, such as on a nightly basis. A manual process could also be employed for cases where immediate service is requested. Enter the MAC address in NWIS, enter it in the DHCP server, and the user is ready.

Another method of documenting MAC addresses could be a Web page where a user can obtain DHCP service by simply entering pertinent information such as their machine name and MAC address. An automated process could then make the necessary database updates mentioned above, and the user is ready. An obvious disadvantage is that a user may not know the MAC address or how to find it. However, the method of determining the MAC address could be described at the Web site.

To make DHCP service more transparent for new users, newly leased machines could come from the distributor or CSU configured to use DHCP. After registering the machine in NWIS, the user simply plugs in and begins communicating.

Mobile Wireless

Another method of enabling mobile computing is to utilize wireless technology. This technology in its basic form is simply a wireless extension to a typical wired network. For example, where a computer traditionally is wired into a wall jack to access the network, a wireless connection would allow this same computer to move about the area serviced by the wireless base station, typically called an access point. This gives the computer user mobility within the range of the wireless access point, typically tens to hundreds of feet within a building. Additional access points increase the range of the wireless network.

An implementation of wireless networking could involve intra-building as well as campus-wide communications. Wireless within a building poses the biggest challenge due to walls and ceilings that cause many RF reflections and the composition of walls which can greatly attenuate an RF signal. Putting aside these sticky points, wireless networking can give a nomadic worker total freedom within range of the network signal. Throw away the wire that ties us down. Use the computer in the office, then down the hall in the conference room, then between buildings when someone stops you and asks about some data as you are headed to a meeting. One could even compute in the bathroom. This would give the busy executive a chance to catch up on e-mail without being disturbed!

Let's start with an example of intra-building wireless. We want to establish a wireless extension of an existing network on the first floor of a multi-story building. Access points would have to be installed in strategic locations throughout the first floor. The number of stations would depend on the size of the area to be covered as well as the type and number of walls. Metal walls will considerably attenuate the RF signal and therefore increase the number of needed base stations. Manufacturers will provide guidelines for establishing the number and location of base stations, but unfortunately, experimentation is necessary. Each base station must be wired into the network. Each computer that is to be mobile must have a transceiver installed, along with an appropriate software driver. Transceivers are typically a PCMCIA card that simply plugs into a laptop, with a small antenna of three or four inches, extending up from the PCMCIA card.

With a client machine configured and access points operating, a mobile computer is ready to communicate. All communications are transparent to the user. The user is automatically connected to whatever network the access point is connected to. This brings up the issue of roaming. Our mobile worker can transparently move about the first floor because the network connection will be automatically handed off from access point to access point just as a cell phone is automatically transferred from one cell to another. In addition, access points can be overlapped in order to increase aggregate throughput.

Will the wireless network operate on the second floor? It may or may not. This depends on the structure of the building. If it does not, but should, additional access points must be installed on the second floor.

Now let's extend this network to another building for which there are no available cables to wire the buildings together. These building could be up to several miles apart. Simply install a pair of wireless bridges with high-gain antennas at some high point on each building. The bridges are wired to the each building's respective network. We now have network connectivity between two buildings that would otherwise not have any ready means to interconnect. Remember, however, that bandwidth is typically limited to 1 to 10 Mbps.

Virtual Desktop

A different approach to mobile computing is to keep the computer stationary and only move the user. The concept of the virtual desktop is that a user's computer screen configuration can be duplicated on another computer, appearing as if the user is sitting at his or her own computer. This could be applied to a conference room that is set up with computers that are available for meeting participants. A user could log on to the computer, see his desktop configuration, and have access to his files. Another user subsequently logging on would see her configuration and files and be able to run her own set of applications.

This project anticipated investigating virtual desktop technology with the release of Windows NT 5.0. However, this software release was delayed such that it was decided to defer this work to the future.

Authentication

Project Scope

This project ostensibly is concerned with making Sandia computer resources available to professionals when they are away from their desks. The boundary for this availability was never precisely defined because of time constraints and philosophical considerations. However, it seemed to be a mistake to implement a system capable of serving mobile professionals in all of Sandia's Tech Areas and ignoring the broader issues. People want to access their resources from home and from hotel rooms around the world, as well as from the conference room down the hall. Authentication plays a role in all these scenarios, with particular emphasis on remote access. Emerging authentication technology was investigated to determine applicability to mobility that goes well beyond the current SecurID system in use at Sandia. Authentication using Public Key technology is desirable due to adoption of PKI by many authentication vendors and an already established PKI infrastructure at the DOE laboratories. New types of tokens show promise for simplification of authentication as well as additional features. There is a growing interest in accessing corporate networks from the Internet, and data encryption or Virtual Private Network (VPN, also known as an encrypted tunnel) is necessary in such a circumstance.

The following sections describe emerging authentication technologies that are potential improvements to existing authentication, including the desire to utilize PKI, the status of the authentication marketplace, the latest in use of tokens, VPNs, and the software selection process utilized in this project.

Requirements

The choice of authentication system hinged on a number of requirements. As usual, many of these requirements would be more ambitious if the market could support them.

- The system must be robust.

- As a pilot, the system should embrace new technologies and provide information on the maturity of these technologies, the level of security provided, their usability, and their general acceptability in our environment.

- If the system uses a public key infrastructure, it must be compatible with the Entrust infrastructure currently in place in the major DOE labs.

- All communications on wireless links and over the Internet must be encrypted. An encrypted tunnel is preferred over encrypted packets[d].

- It must provide VPN capabilities, including encrypted tunnels. Tunnels are needed to protect our data on wireless links, on the Internet, and in many cases to protect sensitive data within our own networks. Other features of VPNs, such as single sign-on and authorization capabilities, are needed at Sandia and would be extremely useful in this project to enhance usability. This need is so great, and the capabilities are being addressed by so many vendors, that the lack of these features is sufficient to disqualify a system from consideration in the selection process. In addition, there is need for backward compatibility with legacy systems such as Kerberos, as well as the capability to work with large applications such as database systems and administrative information systems.

- The system must provide strong authentication. It should address current threats and should deal with known attacks to the extent possible.

- It must allow two-factor authentication; something you know (a password) and something you have (a token). It must also be more convenient than the SecurID cards now in use.

- All users requesting authentication from outside our network must possess a token.

Why PKI?

Sandia has a Public Key Infrastructure (PKI) in place. Authentication using public key techniques is one of the most secure methods available. It does not fall to replay attacks (the process of intercepting a password and using it), it does not involve transmitting passwords, and the authenticator is quite immune to cracking or to attacks involving guessing. In addition, commercial off-the-shelf solutions for authentication using PKI are moving toward support of single-sign-on, encrypted tunnels, multi-client support, control of authorization, and other desirable features. Finally, the vendors see the opportunity to really make the authentication system immune to attacks that have traditionally been extremely difficult to thwart. These attacks have been popular from the time interactive systems first became available, and they will finally be defeated. For these and other reasons, Sandia has been hoping for an opportunity to embrace this technology for controlling access to our restricted network.

The Marketplace

Vendors are just now able to use new groups of techniques to field flexible products that address a significant portion of the authentication challenge that Sandia needs to attack. New capabilities

[d] An encrypted tunnel encrypts the entire original packet, both header and data, and encapsulates this inside another packet. An encrypted packet only encrypts the data, with the original header left intact. Thus an encrypted tunnel provides better protection since the source and destination addresses are hidden by encryption.

are being developed as fast as vendors can bring them to market, and every vendor is advertising their plan to market a much more complete authentication system than they can now provide. In fact, distinguishing software from vaporware is a challenge right now. The same requirements might dictate a different system a few months hence, but the marketplace will evolve rapidly for some time to come.

Systems suitable for this project are emerging. Encrypted tunnels have been around for some time but the systems have lacked flexibility; some connected client and server, others worked only between servers, and a few could connect two clients. In addition, many systems now are limited to a PC/Windows NT environment. Often the single sign-on features really only provide automatic authentication to the Windows NT network after passing through a gateway, and it is only occasionally possible to work with Unix systems or to address legacy applications.

While many vendors work in a public-key environment, only a very few really understand what a public key infrastructure entails. Several vendors want to be suppliers of the Certification Authority software including key management and distribution, as well as the capability for implementing authentication and authorization using this technology. These emerging public-key systems fall far short of our scope because they are incomplete. For example, no known PKI systems are able to deal in a general way with keys from other Certification Authorities, and only a very few are able to do any verification of the validity of the key being used. Sandia's Certification Authority performs many of these functions, and these types of capabilities are regarded by Sandia as requirements.

Choosing a Token

The selection of a token is an important part of authentication system selection. The user must carry the token to access the system – form factor is crucial. For the average user, perhaps half of the interaction with the system is dealing with the token – convenience is crucial. The features contained in the token dictate the level of security available to someone completely outside our network, and also the immunity to attacks from insiders and Internet crackers. Finally, the flexibility of the technology upon which the token is based is the indicator of its ability to deal with future attacks.

The smart card token has the most desirable form factor available today. It is the size and thickness of a credit card. It does not break when dropped or flexed slightly. This type of card is much easier to carry than the rigid, breakable cards in use today. In addition, some vendors have other form factors available, such as the STU-III key form factor. Some Sandians already carry a STU-III key and many of them like that form factor. An informal survey of 20-30 Sandians indicates a general trend of preferences. While there are a number of exceptions, among those surveyed (staff, secretaries, and managers) the most common preference among females was the card ("I already carry several credit cards. What's another one?"). And males tended to prefer the key ("I carry too many credit cards, and a key will hang on my badge.").

Some smart cards can be inserted into a carrier that plugs into a PCMCIA slot. This may be the desirable interface in some cases. The alternative is an external reader the size of a deck of playing cards with a cable to plug into the serial port. Many people object to carrying a separate device and dealing with cables. Having a choice available is important if one type of port is already monopolized. *(also, there may be an issue with the inability of the client computer to support another serial device as all ports may already be in use).*

The convenience offered by a public-key-based smart card seems to be the best available for strong authentication. While many companies are still working in that area, the public-key technology has the capability to accept a password and then to transparently satisfy an unlimited number of requests for authenticators without further intervention by the user. It also can provide link encryption or bulk encryption as needed, without the need for key negotiation.

The process of selecting an encryption key is probably more likely to introduce security holes than any other task the system must perform. The use of public key technology moves that task into a specialized security system designed specifically for that purpose.

The PKI paradigm is robust. It offers good immunity to replay, guessing, cracking, and other current threats. In addition, the configuration of the card allows security to be augmented, usually without altering the user interface (changes may occur that are invisible to the user).

Some cards have only a memory function. This keeps the user's keys on the card rather than on a disk that might be accessed whenever the machine is attached to the network. A more secure version of a smart card not only holds the keys but also performs encryption and signing functions. This type of card does not recognize functions to export keys outside the card, and to subvert the system a cracker must communicate with the card from the user's computer. A higher level of security involves a PIN keypad built into the smart card reader. The PIN travels directly from the keypad to the card, and never enters the computer in any way, at any time. Keyboards with built-in readers achieve part of this goal, but fall short when the user receives a bogus prompt to enter the PIN, instigated by a hacker running code on the user's computer.

Finally, the better smart cards contain chips that are computers in their own right. This provides a very high level of flexibility to respond to future threats. The PKI technology can tolerate changes in implementation strategy, primarily because it is built on a very old mathematical puzzle; the factoring of large numbers. Euler[e] himself is known to have spent a great deal of time on this problem. With this sort of simple and impenetrable base, many implementation details can be modified without fear of providing clues for cracking keys.

A survey of the market at the RSA Data Security Conference (a conference sponsored by RSA Data Security, Inc.) found a lot of similarity in smart cards but not much compatibility. First, there is a standard form factor, including placement of the contacts, size, thickness, and other physical parameters. One vendor also has the STU-III key form factor. Most of the cards have similar functions, at least within the groups outlined above, and of course, some have more memory or better processors. Some vendors have better interfaces, such as the carrier for the

[e] Leonhard Euler, the leading mathematician of the 18th Century.

card to allow insertion in a PCMCIA slot. The best overall set of features seems to be from DataKey, Inc., which has both of the innovations mentioned above. They also have fast on-card processing and a convincing story about cryptographic precautions such as salting of PINs and good random number generators. Unfortunately, there is little effort among the vendors to provide compatibility aside from the standard form factor, and software systems accept only cards envisioned by the implementers.

Although the DataKey smart card was the card of choice, the software selected (see below) accepted only cards from Schlumberger, Limited. Giving up the DataKey STU-III form factor was disappointing, but did not seem to be a major compromise. In addition, the implementation could interact only with memory on the smart card. This is a real security issue, but in a pilot system it seemed acceptable since the option existed to develop software to accept full-function cards. Conclusions from the pilot project would still be valid since the human interface would not change if the cards were upgraded – that change would be invisible to the user aside from performance considerations in a few exceptional cases.

In the longer term, external factors suggestedconsideration of software enhancements to embrace both the full-function cards and pin-pad readers for the pilot project. While this was a feasible goal from the technical viewpoint, the time needed for software development was not compatible with the project timelines.

Selection of Software

Authentication software was surveyed at the RSA conference, including KyberPASS, CyberSafe, SnareWorks, Entrust (a developing partnership with another vendor) and others. After talking with colleagues, some of the possibilities were dropped and more inquiry was performed.

The RSA conference revealed that several virtual network providers address connections between servers only, and some vendors are not equipped to do adequate authentication. Two vendors seemed to address the needs of this project.

Presentations were given by these two potential vendors of security software. Entrust presented a solution that builds on Sandia's current PKI infrastructure and which would meet authentication requirements and provide privacy primarily for non-sensitive data. SnareWorks presented a solution that would meet authentication requirements and combines the advantages of functionality from the PKI infrastructure with providing some backward compatibility into the Kerberos and DCE world. In addition, SnareWorks promises a more practical solution, providing much more flexibility and better services to the customer, including privacy for sensitive data. The cost and effort involved in these solutions seems to be commensurate with their relative value.

Entrust

The Entrust solution provides a means for someone to authenticate to a server (which could be a gateway) and to be logged into a network domain (of Windows NT machines). An encrypted tunnel is established between the client and the gateway, but there seems to be no provision for protection of data between the gateway and server. There is no promise of further services, and there seems to be no way to distinguish between people who are mobile inside the IRN (who would not need a hardware token unless wireless communication were involved) and those who are coming in from the EON (who would always need a hardware token). This situation shows two prominent disadvantages:

- Either everyone would need hardware tokens, or there would have to be two separate servers to handle the mobile worker inside the IRN and those outside the IRN.

- This sort of service is great for technically savvy users, who know how to navigate the network and don't mind logging into their desktop machine separately (this might turn out to be a non-issue for Windows NT 5.0 machines). But the typical manager/executive doesn't want to type more than one password if possible (this can't really be done right now) and they certainly don't want to have to locate their desktop machine inside the Sandia network and then log on. Moreover, some large fraction of these people will not be happy with a solution which shows a different set of icons in a different place on their screen, or which requires a slightly different command to access a frequently-used service. These people will demand training and/or hand-holding which implies huge deployment costs. In short, this might demonstrate the concept, but it doesn't seem to be something people would be happy with, nor does it address the range of services needed for a practical solution.

On the other hand, deployment of the Entrust solution for mobile computing would be relatively simple and painless, aside from dealing with beta software. This sort of effort might be a couple of months for a couple of people. Software is available for evaluation at little or no cost, and tokens must be purchased.

SnareWorks

SnareWorks promises to provide encrypted tunnels from the client to the gateway, and from a gateway to the required service. Authentication at the gateway can be dependent on the location of the user (thus knowing who needs a token). It allows the user to automatically be logged into a pre-defined resource, and tasks launched on that resource. In addition, Kerberos tokens can be obtained automatically and DCE services are available, including authorization. In short, it promises power and convenience for the customer, and backward compatibility. It too is beta software in some sense.

On the other hand, procurement and installation of SnareWorks is likely to consume significant resources. Up-front software cost would be around $15K, which includes a week of on-site analyst support for installation and customization. Some coding may be needed. The magnitude of this effort is unknown. Best guess of the effort is on the order of two months using one and a half people, which would result in a decent test platform with more work needed for actual deployment. This includes some expected tweaking and coding, and it may be cost-effective to contract with the vendor for some of this support. In addition, tokens and other items would be needed.

SnareWorks Chosen for Authentication

A remaining question to be answered regarded the additional burden that might be encountered with SnareWorks's DCE infrastructure. Sandia's DCE technical people suggested that the configuration and management burden of DCE would be less than originally envisioned and that the benefits were substantial. Therefore, SnareWorks was chosen over the Entrust solution.

In making this choice, a less desirable smart card was accepted, from Schlumberger, which is strictly a memory card. This is adequate for this pilot project, but would be marginal for a production application. However, Intellisoft Corporation, developers of SnareWorks, have offered to add support for other smart cards. Given the opportunity, a card with on-board encryption and PIN-pad readers would be preferred. It is not clear that all of these hardware features, plus compatibility with Entrust and an acceptable software system, is available.

This choice of smart card was re-visited after external considerations enhanced the need for testing systems with better security. Intellisoft's offer to develop software was considered, but rejected because it would have unacceptably delayed completion of the pilot. In the rejected scenario, the final configuration would have included full-function smart cards as well as readers with integral PIN-entry pads. This would maximize the security of the system within limits that are both convenient and currently achievable.

In the near future, it is expected that switching to different cards will create little impact. A new standard interface for accessing smart cards is being accepted by vendors, called PC/SC. In a year, this should begin to solve the compatibility issue.

Product Testing

SnareWorks

SnareWorks is a security framework for TCP/IP and legacy applications. It provides authentication with single sign-on capability, standardized authorization across all applications, data protection services, dynamic support for new protocols through Protocol Support Modules, auditing, centralized management, connection monitoring, virtual private network services (VPN), public key certificate services, a downloadable client with zero management, and Web-only clients.

From a SnareWorks glossy: "*SnareWorks is a family of security products that combine the features of private key and public key security technologies to provide TCP/IP-based commercial off the shelf (COTS) and legacy applications with a comprehensive security framework, that includes secure single sign-on, in a completely transparent manner. SnareWorks is a unique solution: it provides any TCP/IP application with strong encryption, a variety of means of authentication, a scaleable authorization model and single sign-on all without requiring any application modifications...*"

SnareWorks was set up to test its ability to provide authentication, encryption, and access control. There are two methods for accessing the IRN using SnareWorks. One method is dial-up through the Public Switched Telephone Network (PSTN). The other method is to connect from the Internet. Authentication was conducted with dial-up and with network access from both outside and inside the Sandia firewall. Encryption was also utilized in all three situations. Access controls were established to demonstrate the ability to finely tune object access such as a particular web page or a particular disk share.

Dial Access

Authentication over a dial-up line is a mandatory function, while access control has merit. However, encryption over a dial-up line is an area of controversy. There has been some discussion about the need to encrypt in this situation. The current production system does not encrypt data. However, with the ever-increasing complexity of the PSTN and ever-increasing hacker activity, one can perhaps no longer be completely sure that data traversing the PSTN is inherently secure. Therefore, dial-up encryption capability was tested.

In utilizing SnareWorks for authentication, encryption, and access control over a dial-up line, the client computer must have the SnareWorks client software installed. The installation is a simple matter of executing a setup program from floppy disk, CD-ROM, or from the network. Once installed, the client must have a gateway definition, which in this case is the SnareWorks server, located at 134.253.4.53. This gateway is configured so that the target is all addresses in the IRN; 134.253.0.0. With this setting, whenever the client attempts a connection to any target in the IRN, as shown in Figure 10, the connection is first established to the SnareWorks gateway in an encrypted mode, then in unencrypted mode to the target machine. Therefore, all data from the client to the gateway in the IRN is encrypted and protected from eavesdropping anywhere within the PSTN.

Testing dial-up encryption showed that this technology is functional. Encryption puts a heavy strain on the client CPU and reduces throughput, but data was successfully passed between the client and any selected targets.

Due to time constraints and limited hardware access, dial-up authentication was not tested. However, the concept is simple as shown in Figure 10. Whereas with the production system that currently authenticates with SecurID, a change would be made to the dial-up equipment to hook into SnareWorks. When SnareWorks gives the OK on authentication (i.e.: the client smart card is opened by the user and the correct data is passed to SnareWorks), the dial-up connection is established.

Access control is a complex subject due to literally an infinite number of possible rules that can be established. With respect to dial-up, SnareWorks can be configured to control access based on the location of the client, granting or denying access to certain resources when the client is dialing in from the outside world. For example, dial-up users might be restricted to a particular set of resources, or denied access to certain resources. This function was tested by establishing rules governing access to a test Web server. Certain pages were set for all access, some for read only (GET in the web vernacular), and some for write only (POST in the web vernacular). After dialing in with encryption (but being authenticated using SecurID on the production system), the access controls worked as expected.

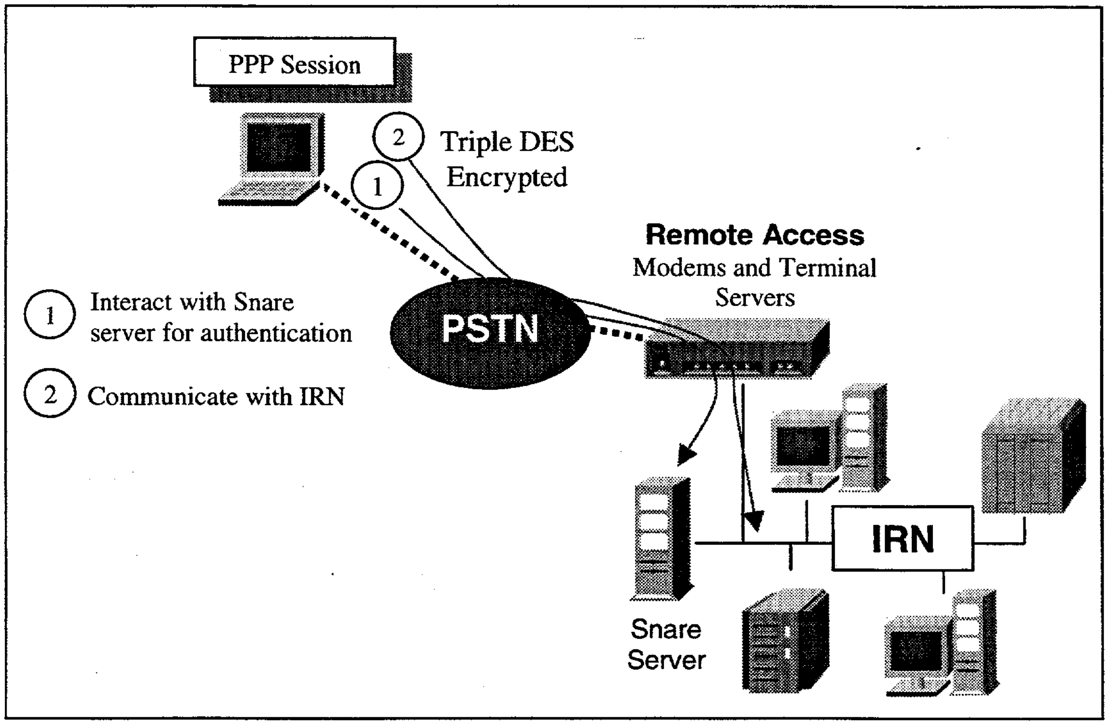

Figure 10. SnareWorks Dial-Up Access to the IRN

Internet Access

An expanded approach to accessing the IRN is to connect from the Internet. This implies dial-up, with the distinction that access into the IRN is by way of dialing into the Internet, rather than dialing directly into the IRN. Internet access is possible only as a Telnet session with the existing production system. However, using SnareWorks with encryption, this can be expanded to include virtually any connection mode. Someone at home or on the road that has already connected to the Internet through an ISP may prefer to access the IRN directly rather than having to disconnect and dial into the Sandia modems.

IRN access through the Internet can be accomplished using SnareWorks in a manner similar to direct dial-up, the main difference being the need to penetrate the firewall that protects the IRN from the Internet, as shown in Figure 11. The firewall must be configured to pass data on port 509. The client again must be running the SnareWorks client software. The client software will wrap all packets into port 509 requests that pass through the firewall to the SnareWorks server.

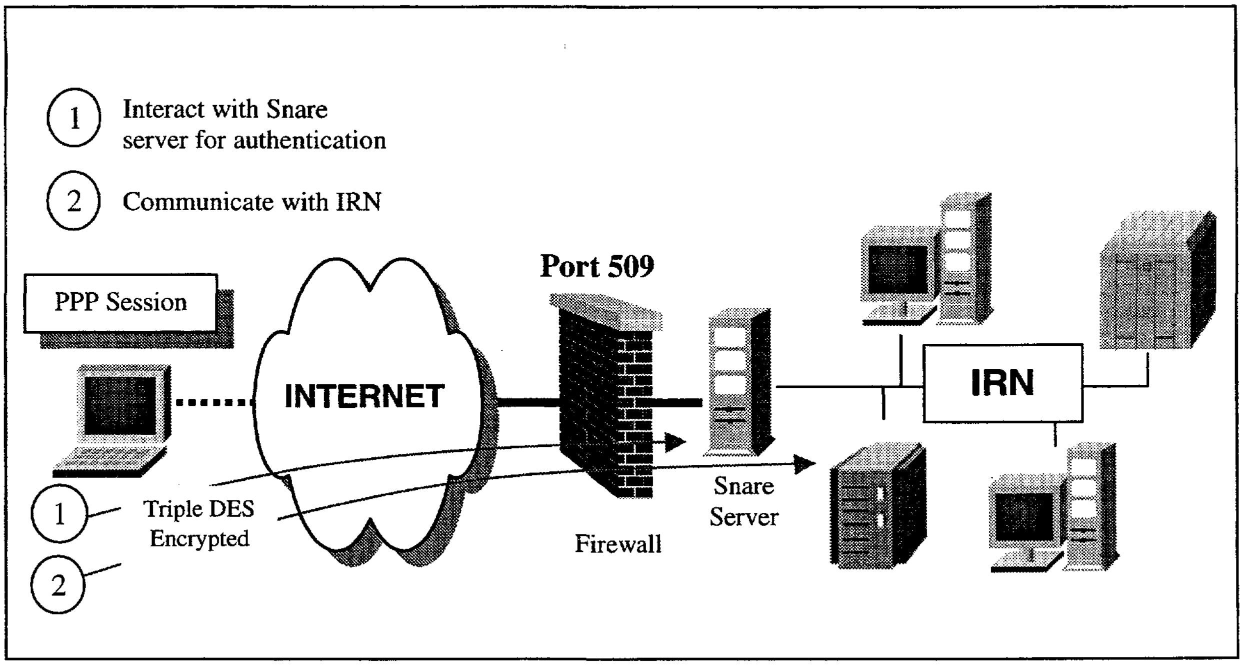

Figure 11. SnareWorks Internet Access to the IRN

The server then unwraps the packets and they go to their intended destination. The client software must be configured to use the SnareWorks gateway for accessing any addresses within the IRN (134.253.0.0). With this configuration, a user could dial into an ISP and then request some type of connection to the IRN such as web browsing. The SnareWorks client will detect a connection request to an IRN address and initiate authentication. The user will be prompted to enter a PIN that enables their smart card as shown in Figure 12. If the PIN is correct, the card is enabled and the SnareWorks server processes a unique, single-use authentication data packet created and digitally signed by the card.

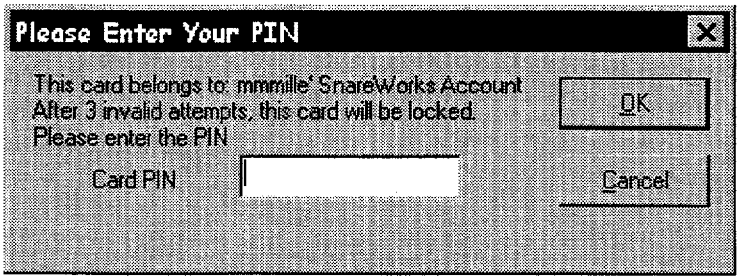

Please Enter Your PIN

This card belongs to: mmmille' SnareWorks Account
After 3 invalid attempts, this card will be locked.
Please enter the PIN

Card PIN

Figure 12. SnareWorks Client PIN Prompt

If the authentication is successful, the requested connection is completed, and the user receives a confirmation message as shown in Figure 13. All subsequent connection requests prior to a disconnect are completed transparently to the user using the smart card authentication data.

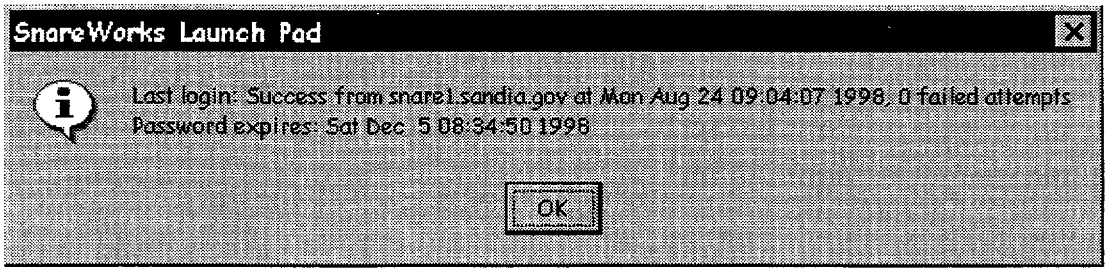

SnareWorks Launch Pad

Last login: Success from snare1.sandia.gov at Mon Aug 24 09:04:07 1998, 0 failed attempts
Password expires: Sat Dec 5 08:34:50 1998

Figure 13. SnareWorks Login Confirmation

When using a Web browser in this scenario, proxy settings should be turned off. This allows traffic destined only for the IRN to be encrypted, and all traffic destined for the Internet to pass unencrypted directly to Internet destinations. An example would be when working from home, a connection is established to the IRN to update a timecard, while simultaneously surfing the Web for sports news. With a proxy setting on, ALL data will pass through the encrypted pipe to the IRN, including non-Sandia traffic destined for the Internet. This works but is not at all ideal, as the Sandia proxy server must handle all out-bound Web traffic, which should be going directly to its destination. Rather, with the browser proxy turned off, all non-Sandia Web connections will go directly to the intended destinations and Sandia traffic (both Web and otherwise) will go only to Sandia computing resources.

The concept of IRN access from the Internet can be taken even further such that Sandia maintains no dial-up service other than authentication. Leave the modem work to an ISP – that is their job. Sandia could contract with a local ISP or a national ISP, which would provide local dial-up from most locations, saving expensive long distance and toll-service charges. One such service is called iPass, from iPass, Inc. With this service, a corporate employee on travel can make a local call in major cities throughout the world to connect to the Internet and the company's intranet. Data protection is provided by VPN by either iPass or the customer's corporate authentication service.

An additional application of Internet access is the potential to eliminate at least some leased-line services. In cases where a leased line is in place, but high-speed access is not critical, the Internet can be leveraged as a lower cost path into an intranet in combination with VPN technology.

One obvious down side to these approaches is the reliance on a contractor, putting direct control of remote access out of Sandia's hands. Nevertheless, this might well be overshadowed by the benefits gained from not having to maintain dial-up hardware and from saving costs by eliminating long distance and toll-free service charges.

Internal Access

SnareWorks can be used within the IRN for access control and encryption. Authentication can certainly be applied, and has the added benefit of single sign-on whereby a client, once authenticated, can bypass further authentication interactions by having the SnareWorks client automatically respond to authentication requests from other applications. Encryption can be mandated to specific IRN services by running SnareWorks on the server that is providing each service. Any user that wishes to access this information must be running the SnareWorks client in order to access encrypted data. Very flexible and comprehensive access controls can be placed on desired services in the same manner as providing encryption.

SnareWorks Comparison to SecurID

The authentication tool in production at Sandia is the SecurID token from Security Dynamics Technology, Inc. A non-predictable, one-time-only access code is generated by a SecurID card every sixty seconds in this relatively simple system. The user is authenticated using this access code in conjunction with an authentication server and a PIN known only to the user. The SecurID card is a credit card-sized device with a metal case to protect the electronics, making it considerably thicker and more rigid than a credit card.

The following table compares features of SecurID and SnareWorks:

SecurID	SnareWorks
Simple system – if the user is successfully authenticated, the door is opened to the protected resources, with all access then controlled by those resources	Complex System – authorization, access control, and encryption
Only a credit-card sized device is needed with no physical connection to anything	In addition to a credit-card sized device (smart card), a reader is required which must be plugged into the computer (requiring an available serial port), until such time as a PCMCIA card carrier is available for laptops
Although only the size of a credit card, the SecurID card is rigid and subject to damage if sat upon or dropped	The smart card is lightweight, flexible and not easily damaged
Will interact with any OS since it is simply a passcode generator	Client software must be installed on user computer, limiting the choices of operating system
Can be clumsy to use when entering PIN into card; user must read and type the generated passcode	No need to enter PIN into a small card; all interaction is either on the computer screen or this is no interacton
Authorizaton control only	VPN capability, authorization control, and access control
Timing is important when entering the generated passcode in response to a login prompt – the passcode may expire before the authentication takes place, requiring the user to enter the PIN again to generate a new passcode and the card can drift out of synchronization if not used regularly	No synchronization issues

Table 1. Comparison of SecurID and SnareWorks Features

There are two methods for accessing the IRN using a SecurID card as shown in Figure 14. One method is to dial-up through the Public Switched Telephone Network (PSTN). Another is to Telnet from the Internet. In both cases, access is controlled by the SecurID server, which authenticates based on an access code that is generated in conjunction with a PIN entered by the user.

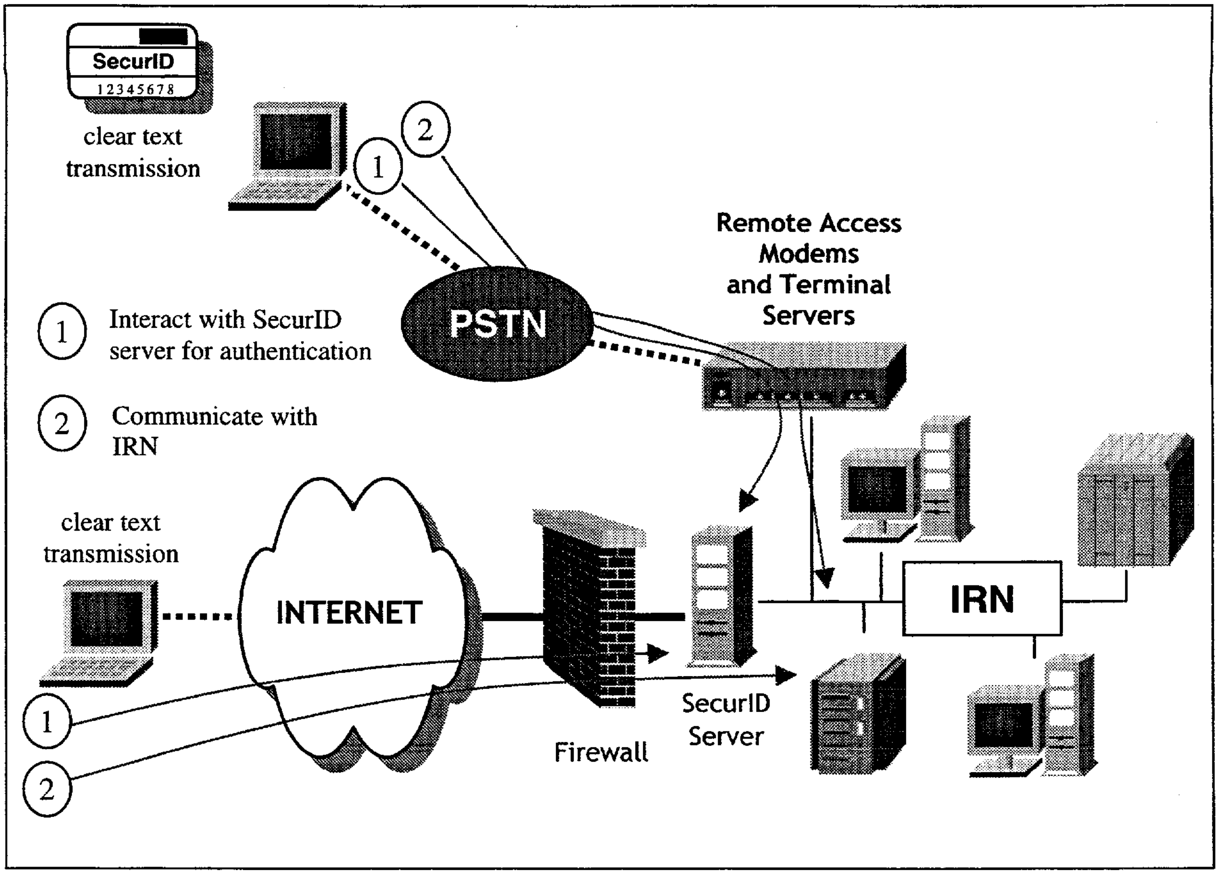

Figure 14. SecurID Access to the IRN

SecurID Authentication

SecurID is a simple authentication methodology. Access is granted to a resource based on a username and passcode combination. The passcode is a six digit number which is generated after the user enters a PIN into a credit-card sized device called a SecurID card as shown in Figure 15.

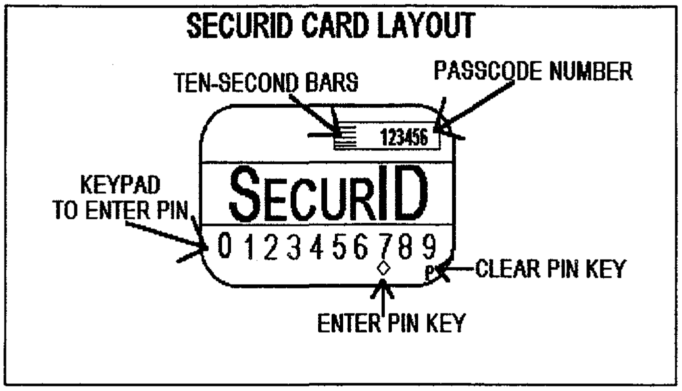

Figure 15. SecurID Card

The SecurID server is time-synchronized with the card. The combination of the user PIN and the card generates a passcode that is verified by the server. By use of a PIN, each user generates a unique passcode at any given time.

Accessing the IRN through Dial-up

A user enters a username and passsword (for Windows logon) and dials a modem bank (Figure 16). When the modem answers, the user is prompted to enter a username and access code (for SecurID) as shown in Figure 17. This prompt is generated by the SecurID server, which will attempt to authenticate the user. The user enters the username, then enters a PIN into the SecurID card. The number generated by the SecurID card is then entered at the access code prompt. If the authentication is accepted, the user is connected to the terminal server. At this point, a command must be entered to initiate a PPP session. After this, the user is connected to the IRN.

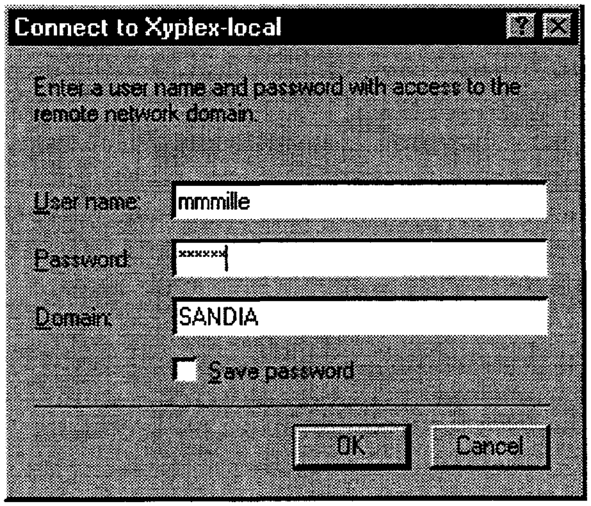

Connect to Xyplex-local

Enter a user name and password with access to the remote network domain.

User name: mmmille

Password: ******

Domain: SANDIA

□ Save password

Figure 16. Prompt for Windows Username and Password

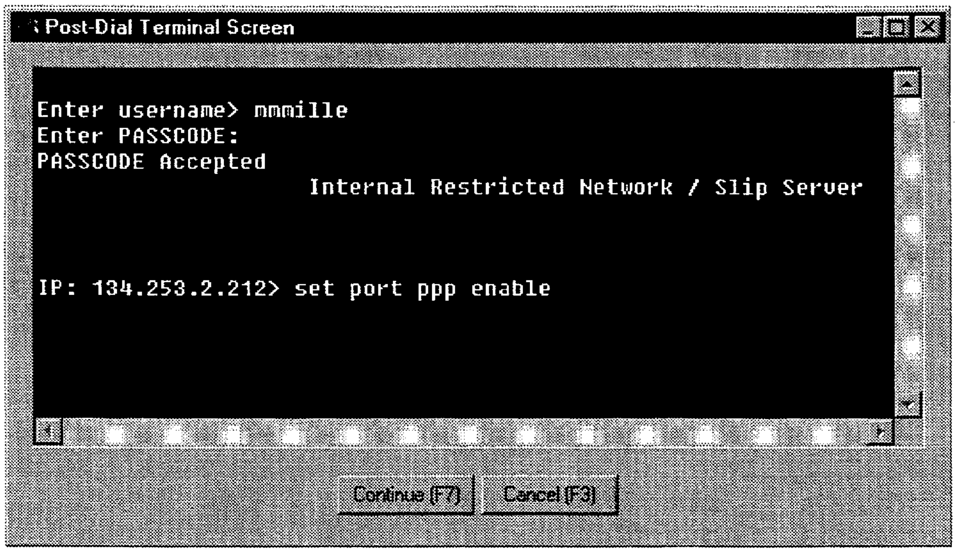

```
Enter username> mmmille
Enter PASSCODE:
PASSCODE Accepted
                    Internal Restricted Network / Slip Server


IP: 134.253.2.212> set port ppp enable
```

Figure 17. Pop-Up Window for Authentication and PPP Session Start

There is an optional method of connection which somewhat automates the process. Using a script, dial-up networking can interact with the SecurID server, sending the username, access code, and PPP command. The user must enter the username and access code at the beginning of the dial-up session. The disadvantage to this method is that timing is important; if the access code generated by the SecurID card times out before the script completes, the user will be rejected. Therefore, the session should be initiated at the beginning of a SecurID time interval. Also, since the SecurID passcode must be entered at the beginning of the session in place of the Windows password, the user is prompted for the Windows username and password after the connection is made.

Accessing the IRN through Telnet

With this method, a user is already on the Internet by some means such as an ISP or other corporate LAN. A telnet session is initiated by connecting to the Sandia SecurID server at IP address 134.253.26.5. The user is prompted for a username and passcode, just as with the dial-up method. If the SecurID server authenticates the user, a connection is established with a terminal server. The user can then telnet to destinations in the IRN.

With both dial-up and Telnet to the IRN, all data is transmitted in clear text, including usernames.

Performance Issues

SnareWorks

Using SecurID, no performance penalty is realized. SecurID is for authentication only, and gets out of the way of transmission once a connection is established.

However, with SnareWorks, encryption places a heavy burden on a client CPU. A 133 Mhz PC was running at 100% CPU when testing file transfers with encryption over a LAN. The CPU must encrypt or decrypt each packet depending on the direction of transfer.

Throughput was reduced by 80% over a non-encrypted transfer. However, the vendor has claimed "significant improvement" in throughput with the latest software release. This had not been tested prior to this report.

Authentication, encryption, and access control can be bypassed when utilizing SnareWorks within an Intranet such as the IRN and using the gateway mode (all connection requests to a resource pass through the SnareWorks gateway). The user simply disables the client software by turning it off or uninstalling it and connects directly to the desired resource. However, having SnareWorks running on a server as well as the client protects the server from client bypass.

The card reader tested for this project is the size of a deck of playing cards with a cable that plugs into both a serial port and the mouse port (in series with the mouse if one is present). It proved to be useable but not convenient. It certainly functioned as intended. However, it adds clutter to the desktop and is yet another device to carry around with a laptop (along with power adapters, communications cables, and external devices). There is also an issue with the need to connect a reader to a serial port. Some machines don't have an available serial port. And in one case, attempting to free a serial port on a laptop to use with the card reader resulted in malfunctioning modem and network cards. A card reader would surely fail to gain much support from typical nomadic workers. Compared to a SecurID card, a card reader is easier to use but this is not likely to make up for its other inconveniences.

A preferred method of card reading is to use a PCMCIA device that either accepts a card or is a card itself. All the aforementioned inconveniences of a card reader are eliminated with this device; no serial port requirement, no cables, and no additional devices to carry assuming the reader is kept in the PCMCIA slot. The primary convenience is simplified operation compared to a SecurID card. The user simply enters a PIN on the computer screen and the authentication software does the rest. Such devices are becoming a reality, however none were available for testing with SnareWorks during this project.

Stability of the SnareWorks server software was a problem. Quite often the server would stop and require a manual restart by an administrator. In working these problems with the vendor, a bug fix would often result. There were also cases where the server software was running but not responding. This would require an administrator to stop and then restart the software. Several bugs were uncovered in the client software as well and were addressed by the vendor as the problems were identified.

Implementation Plan for Providing Authentication Service to Nomadic Worker

An implementation of an authentication service could be constructed as shown in Figure 18. The authentication service centers on the server running SnareWorks. This server is located behind the firewall to take advantage of the firewall's protection. The server must run the Solaris operating system from Sun Microsystems, Inc. A possible server could be a Sun Microsystems Ultra 5 with 64MB of memory and a 4GB hard drive. This server must have DCE installed along with SnareWorks.

There are three access methods that pertain to SnareWorks authentication; dial-up, Internet, and internal. Each access method has its own set of requirements.

As shown in Figure 18, a remote client machine running SnareWorks would dial into the IRN through the telephone system. Authentication takes place by having the modem communicate with a Radius server (an industry-standard authentication server) which in turn communicates with the SnareWorks server. Radius is required since the modem equipment is compatible with Radius and not SnareWorks. The data can be encrypted using Triple DES (encryption of data three times with three different keys).

For Internet access, a hole must be punched through the firewall to allow one and only one service to pass through to the SnareWorks server. Enabling port 509 on the firewall creates the hole. When a SnareWorks client in the Internet attempts to access a service behind the firewall, the connection request is disguised as a Port 509 connection, which passes through the firewall to the SnareWorks server. The SnareWorks server then unwraps the connection request and forwards it to the desired destination. The data is again encrypted using Triple DES. This creates a tunnel through the Internet in that the data is traversing a public network, but is private due to encryption. This is commonly known as a VPN, or Virtual Private Network as shown in Figure 18.

For internal access, where a user is in the IRN, only the SnareWorks server comes into play for any authentication functions. In this case, the SnareWorks client must be configured for gateway mode so that all connection requests that are to be processed by SnareWorks are directed to the SnareWorks server.

Any users wishing to utilize SnareWorks authentication must have the SnareWorks client software installed on their machine. This is a simple process as there is no interaction required on the part of the user. The installation could be automated as is done with SMS application installations for the COE. While there are parameters that can be set within the client, these parameters can be set from the server. This would ensure continuity among clients and relieve users from having to understand configuration settings. A smart card reader must be installed on the client computer for authentication. A typical reader consists of a playing-card sized box with a cable that plugs into both a serial port and the mouse port (in series with the mouse). After a simple configuration process that initializes the card to authenticate using DCE credentials, the client machine is ready. The smart card can contain other authenticating information such as Entrust credentials to be used for authentication with Entrust-capable applications.

The bulk of the effort required to implement a SnareWorks authentication service would involve installing, configuring, and administering the SnareWorks server. First, the DCE services must be installed and configured by a DCE administrator. Once this is complete, the SnareWorks software can be installed by anyone having system privileges. This is a fairly easy process. However, once installed, the SnareWorks server must be configured. This is not an easy process. There are many parameters and configuration options with SnareWorks server. A manager of a SnareWorks server must be well versed in its operation.

Hosts and domains (groups of hosts) must be defined although DCE will provide domains that are already defined within DCE. Machines that are not DCE-enabled will have to be manually added to SnareWorks if they are to be involved in any rules.

Time definitions control when rules apply and are broken down into Business, Weekend, and Everyday. The manager must establish the hours of operation for these time periods.

The connection rules are at the heart of SnareWorks. Here, the manager must establish rules to govern various connection possibilities on a protocol basis. The primary parameters are encryption on or off, hours of operation based on the time definitions, authentication on or off, and auto-login. For example, the HTTP protocol (used for web pages) could be set for authentication for any user at any time with encryption from any machine. With this rule set, a user connecting to the IRN from the Internet would be required to authenticate using a smart card and the session would be encrypted from the user machine to the SnareWorks server inside the IRN.

Object rules can be established to cover access to specific objects such as web pages, shared disks, or ftp sites. An example would be a site in the IRN that is blocked from being accessed from anywhere outside the IRN.

Once all these pieces are in place, a complete authentication service is ready for action. Dial-up is authenticated using smart card technology and optionally encrypted. Network access is available from anywhere in the Internet with encryption using Virtual Private Network technology.

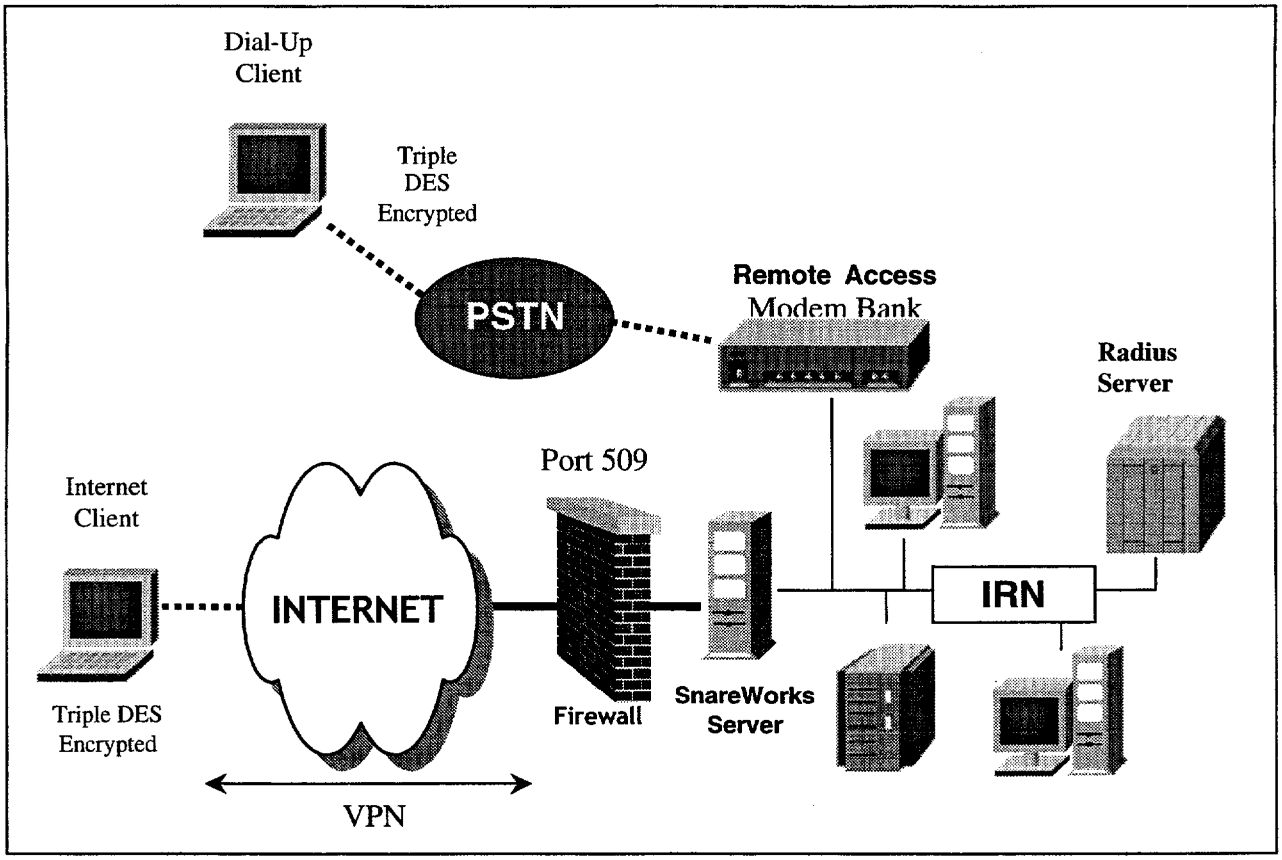

Figure 18. Implementation of an Authentication Service

Conclusion

This pilot study project explored the problem of providing access to the nomadic worker who desires to connect a computer through network access points at a number of different locations within the Sandia/NM campus as well as outside the campus.

This project has shown that location independent computing can be accomplished in various ways to meet various needs. The most significant scenario, mobile wired, can be accomplished using the Internet Standard DHCP protocol to provide dynamic IP addressing to any point within Sandia's IRN network. Dynamic addressing provides transparent network access to the user as the user moves around the network. It has been shown in this report that conference room network access is difficult to obtain. DHCP can be applied to conference rooms to greatly simplify the process of network access. The technology can be scaled to cover the entire corporate network, providing mobility across the enterprise. Additionally, DHCP reduces the administrative burden related to IP address management since addresses are managed automatically, and can be utilized to help maintain the integrity of NWIS data.

Other mobile computing scenarios studied were wireless and virtual desktop technologies. Wireless networking can both provide mobility and enable networking in areas where wire plant is either scarce or unacceptable. Virtual desktop technology could be applied to classrooms, conference rooms, or any situation where a user wishes to access the computer desktop from other locations. In other words, the user is mobile but the computer remains in one location. Testing of these technologies was deferred due to security and time constraints.

Authentication using smart cards and data protection using VPNs was studied and it has been shown that these emerging technologies can be applied to mobile workers outside the corporate intranet to improve the authentication process and provide protected Internet access to the corporate network. For example, the current authentication tool, SecurID, is unpopular among SNL users due to its cumbersome interface. Smart cards can improve on this drawback by transparently authenticating. Outside access to the IRN can be expanded over direct dial-up by leveraging the Internet in conjunction with VPN technology that protects the data from eavesdropping. VPN technology can also be applied to eliminate the need for a corporation to maintain modem pools by utilizing ISP services, the Internet, and encryption.

Recommendations

To achieve mobility within the IRN, it is recommended that the Internet Standard DHCP protocol be utilized for providing dynamic addressing. This is a non-proprietary solution to mobility that can be applied to any part of the Sandia/NM network without regard to particular hardware and vendor implementations. In particular, DHCP can be applied to conference rooms to provide transparent network access, eliminating the need for users to reconfigure their computers or rely on support organizations for network access.

When applied to the enterprise, DHCP reduces the administrative burden related to IP address management since addresses are managed automatically. Additionally, DHCP should be configured to respond only to known MAC addresses. This prevents DHCP from interfering with clients that are expecting service from some other source, and provides a mechanism for helping maintain NWIS data by requiring computers to be registered in NWIS before gaining network access.

Authentication technology needs further study. Products are emerging that can provide improvement over the current SecurID system and provide protected access to the IRN from the Internet, but lack of robustness is an issue. As products emerge that are robust, a prototype system can be established to further investigate applicability to Sandia's authentication needs.

To further enable mobile computing, it is recommended that additional studies be conducted on wireless and virtual desktop technology. Wireless communications are rapidly becoming a viable alternative to wired communications in some situations. A prototype system could be developed to demonstrate mobility; or a wireless service could be employed such as cellular or satellite. With the release of Windows NT 5.0, virtual desktop technology can be investigated and a prototype developed.

References

[1] Excerpt from an e-mail message from J. P. Long, Department 4621, Sandia National Laboratories, November 20, 1997.

[2] R. Droms, *Dynamic Host Configuration Protocol*, RFC2131, March 1997.

[3] J. Postel, Editor, *INTERNET OFFICIAL PROTOCOL STANDARDS, STD: 1*, May 1998.

[4] http://web.syr.edu/~jmwobus/comfaqs/dhcp.faq.html

[5] http://charlotte.acns.nwu.edu/internet/tech/dhcp/ and http://www.net.cmu.edu/services/internet/bootp.html

[6] C. Perkins, *IP Mobility Support*, RFC2002, October 1996.

[7] http://www.cs.purdue.edu/research/crosspoint.html

Appendix A

Evaluation of some DHCP server products and their Potential Application to Network Mobility

Contents

Figures

APPENDIX A

Introduction

Several DHCP/DNS servers have been analyzed for possible application to the Location Independent Professional project. The ability to provide simple network access to a roving staff member can be addressed by using the DHCP (Dynamic Host Control Protocol) protocol to provide network addresses on demand. No prior network knowledge is needed by the user regarding how to configure the PC for network access (other than perhaps configuring the PC to utilize DHCP). In addition to providing a network address with DHCP, automatic update of the DNS (Domain Name Server) server is desirable in order to make the roving node known to other nodes in the network. This is known as Dynamic DNS, or DDNS. This is generally not significant for a client PC, but does come into play for access to some services, such as the news reader news.sandia.gov. This machine will only service clients that are registered in DNS. Logging of DHCP events is very important in order to track occurrences and have the ability to either generate reports on past activity, or retrieve information on specific occurrences. Finally, remote administration would allow a DHCP and DNS environment to be managed from a central site.

Mandatory Features

- DHCP server - a server capable of providing dynamically assigned IP addresses to clients, with a distributed architecture in order to service potentially many network segments (ie: several DHCP servers with a common IP address pool)
- DNS server with dynamic update capability - normal DNS service and the capability to receive name/address updates from the DHCP servers.
- Event logging and auditing - just when was an IP address granted a lease and what machine received it? This is very important for investigating occurrences.
- Remote administration - from a central location, or perhaps any location, the DHCP and DNS servers should be accessible for management purposes with appropriate access control mechanisms.

Some Notes

DHCP Parameter Override

Client network parameters seem to override whatever parameters DHCP attempts to set. For example, when a client has a DNS server defined, this address will be used rather than the address being offered by DHCP.

Node Types

Windows machines are characterized by a node type with respect to network name resolution. These node types define how a machine behaves when attempting to resolve network names using the Windows Internet Name Service, or WINS:

- b-node
  A computer configured as a b-node will use broadcasts to resolve names on a network. If a computer receiving the broadcast on the local network segment recognizes its name, it will respond with its address.

- p-node
  A computer configured as a p-node will use a WINS server to resolve addresses rather than broadcasts.

- m-node
  A computer configured as an m-node will first try a broadcast for name resolution and, if that fails will switch to the p-node mechanism and query a WINS server.

- h-node
  A computer configured as an h-node will first try a WINS server for address resolution. If that fails, it will resort to broadcasts.

The significance of these node types with respect to DHCP is that a client that is configured to accept all parameters from the DHCP server should receive the node type as part of those parameters. The node type h-node is the preferred type, as it will enable a client to query a WINS server directly if one is available, or broadcast for the desired name information if no WINS server is available. If the node type is not set by the DHCP server, the client may or may not have the parameter set and may act in a less than desirable fashion. For example, when set as a b-node, a computer will only broadcast when attempting to resolve a computer name. If no WINS server is available in the local network segment, no name resolution can take place beyond the local network segment. The computer will be unable to communicate with nodes outside the local network segment. However, if set to h-node, a computer will attempt to communicate directly with a known WINS server for name resolution

DHCP Parameters As Set In All Trials

DNS server: 134.253.181.25 and 134.253.16.5
WINS servers: 134.253.181.12 and 134.253.181.26
Router: 134.253.4.254 for Subnet 4 and 134.253.5.254 for Subnet 5
Domain: sandia.gov

Many other parameters can be set. These are the minimum for proper operation of the client.

APPENDIX A

Client Configurations

Windows95

For a Windows95 client, the network parameters can be set as follows (assuming TCP/IP has already been installed):

- Open the Network applet from the control panel or right click on the Network Neighborhood
- Select TCP/IP (for the appropriate adapter; some computers have two or more adapters for LAN and dial access) and click the properties button (or just double click TCP/IP)
- In the TCP/IP window, select the IP Address tab and select "Obtain an IP address automatically"
- Select the WINS Configuration tab and select "Use DHCP for WINS Resolution"
- Select the Gateway tab and remove any gateways (routers) that may be defined so that the "Installed gateways" box is empty
- Select the DNS Configuration and select "Disable DNS"
- OK the Network applet and reboot

Windows NT

- Open the Network applet from the control panel or right click on the Network Neighborhood
- Select the Protocols tab
- Select the TCP/IP protocol and click the properties button (or just double click TCP/IP)
- Select the IP Address tab, then select the appropriate adapter and select "Obtain an IP address from a DHCP server"
- Select the DNS tab and remove any DNS servers that may be entered. To delete, select the desired DNS server address in the "DNS Server Search Order" box and click the Remove button
- Select the WINS Address tab, select the appropriate adapter, and remove the primary and secondary WINS servers, if they exist, by placing the cursor in the far right of each box and using the backspace key to remove the entries
- OK the Network applet and reboot

Product Tests

Bay Networks NetID

Bay Networks,Inc.
http://www.isotro.ca/Products/nav/t_netid.shtml
602-957-9707
800-890-8671

Features
DHCP server
DNS server
Dynamic DNS server
Web-based remote administration
Oracle or Sybase database

Requires Oracle 7.x or Sybase 11.0 to 11.4 database as a separate product. A runtime Oracle 7 Workgroup Server can be purchased for NT or Solaris from Bay Networks.

This product will not be tested since the customer must provide one of the above databases when trying the demo.

MetaInfo Meta IP / Manager

MetaInfo, Inc.
http://www.metainfo.com/
Product Information: info@metainfo.com
Sales: sales@metainfo.com

Features
DHCP server
DNS server
Dynamic DNS server
Event Logging
Simple web-based remote administration
Microsoft Access database Version 8.0

The product seems to be a reasonable DHCP server and DNS server. DHCP clients function generally as expected. DNS service was functional. The dynamic DNS function is designed to update the DNS server from the DHCP server, but this function could not be made to work. Configuration was straightforward except for the dynamic DNS. The web-based administration provides step-by-step configuration, requiring no editing of configuration files.

Some Observations:

- There were problems with a laptop not establishing an IP address on boot; gets address sometime after boot-up.

- Logging is poor. A text file dhcp.leases shows the outstanding leases, but no history. This file is used by Meta IP/Manager to show the leases in a Java window and is not intended for direct human viewing. A text file dhcp.log shows server activity, such as what MAC address requested an IP address. But does not have time stamps. This activity can be also logged in the system Event Log, with time stamps.
- No release mechanism for leases. For a lease that is known to be no longer needed, there doesn't seem to be a mechanism for deleting the lease at the server. Leases that have not been used for some time stay assigned, but expired. Perhaps if the expired lease is eventually needed, it would be reassigned.
- Has a Web-based management interface. Allows managing DHCP leases and DNS configuration from any machine that is allowed access. Works well, but would only work with Netscape 4.0.
- Unable to make dynamic DNS function. Probably just a configuration problem, but several days of effort were fruitless.

Sample Leases File

- This file, as shown in Figure A – 1, is used by a Java applet to display the leases. It is a text file, but not designed to be viewed directly.

```
› lease 134.253.4.181 {
› starts 5 1998/1/23 22:00:12;
› ends 1 1998/1/26 22:00:12;
› hardware ethernet 0:60:97:85:c8:3;
› uid 1:0:60:97:85:c8:3;
› }
› lease 134.253.4.185 {
› starts 3 1998/1/21 19:28:06;
› ends 1 1998/1/26 15:02:07;
› hardware ethernet 0:20:2b:3:c:dc;
› uid 1:0:20:2b:3:c:dc;
› host SAGW044;
```

Figure A - 1. Sample Leases File

Sample Log File

- This is a simple text file viewable with a utility such as Notepad, as shown in Figure A - 2.

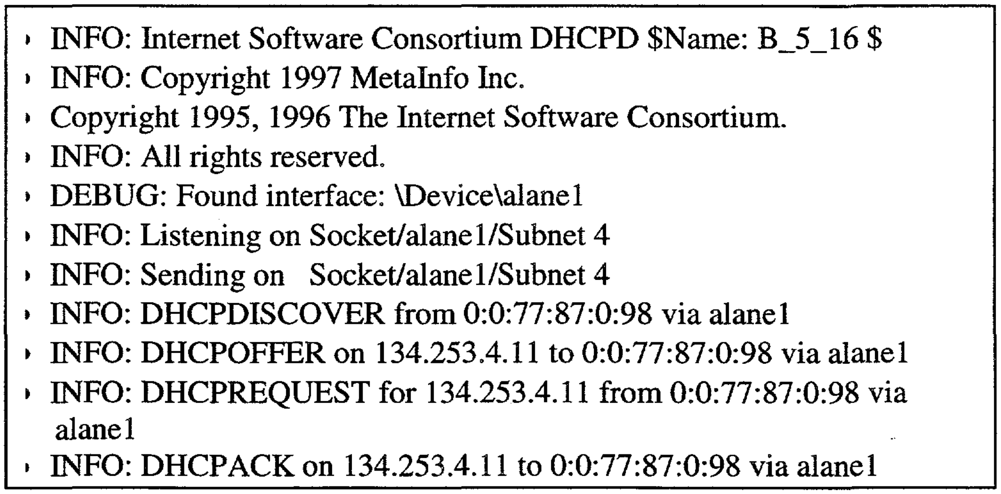

```
› INFO: Internet Software Consortium DHCPD $Name: B_5_16 $
› INFO: Copyright 1997 MetaInfo Inc.
› Copyright 1995, 1996 The Internet Software Consortium.
› INFO: All rights reserved.
› DEBUG: Found interface: \Device\alane1
› INFO: Listening on Socket/alane1/Subnet 4
› INFO: Sending on   Socket/alane1/Subnet 4
› INFO: DHCPDISCOVER from 0:0:77:87:0:98 via alane1
› INFO: DHCPOFFER on 134.253.4.11 to 0:0:77:87:0:98 via alane1
› INFO: DHCPREQUEST for 134.253.4.11 from 0:0:77:87:0:98 via
  alane1
› INFO: DHCPACK on 134.253.4.11 to 0:0:77:87:0:98 via alane1
```

Figure A - 2. Sample Log File

QIP

Quadritek Systems, Inc.
http://www.quadritek.com/
E-mail: sales@quadritek.com
Sales: 888.683.2254

<u>Features</u>
DHCP server
DNS server
Dynamic DNS server
Web-based and native client remote administration
Event logging
Sybase or Oracle database
Auditing
Reporting
A demo version of this product has been tested. The demo is an almost fully functional version. It is limited to 100 IP addresses and does not support remote servers and remote administration.

This very feature-rich product translates into "difficult to set up" but works well.

APPENDIX A

A laptop experiences failure to find a domain server when logging in immediately after a boot. One or two more logon attempts will succeed. This is probably due to the time required to get DHCP to complete. A workaround is to wait some period before attempting the logon. This period appears to be around 10 seconds.

The DHCP server was configured with two "permanent" addresses in order for the corporate DNS server to recognize addresses while still exercising the DHCP function. A block of addresses are set up as dynamic so they can be leased as necessary.

Routing of DHCP Requests

A second network segment was configured in order to test the ability of a router to pass the client DHCP request on to a server in a different network segment. In QIP, a new network segment called Subnet 5 was created and one dynamic address added at 134.253.5.180. Subnet 5 was also added to the list of network segments that SAGW1140 would service. The machine SAGW040 was configured to use network segment 5. The router that services network segments 4 and 5 was modified to forward bootp requests from network segment 5 directly to 134.253.4.35, the DHCP server. This test worked successfully. SAGW040 was properly assigned a Subnet 5 address, 134.253.5.180, and the default router 134.253.5.254.

Auditing

QIP supplies an audit mechanism for showing the change history of items and administrators. For example, given a particular node by name or IP address, the audit will show all changes made to that item. Figure A - 3 shows that IP address 134.253.4.11 has been modified to change the object name from SAGW040 to SAGW044.

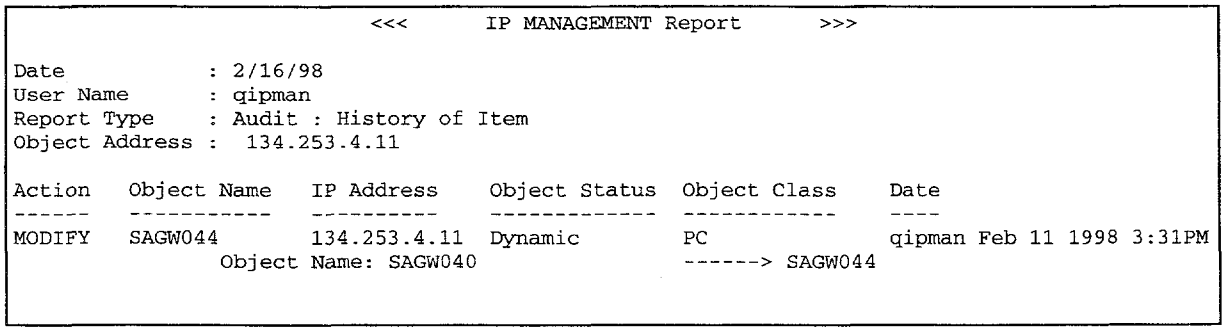

```
                           <<<       IP MANAGEMENT Report       >>>

Date            : 2/16/98
User Name       : qipman
Report Type     : Audit : History of Item
Object Address  :  134.253.4.11

Action    Object Name    IP Address     Object Status  Object Class    Date
------    -----------    ----------     -------------  ------------    ----
MODIFY    SAGW044        134.253.4.11   Dynamic        PC              qipman Feb 11 1998 3:31PM
                 Object Name: SAGW040                  ------> SAGW044
```

Figure A - 3. IP Management Report

Reports

QIP supplies several reports that can be viewed on screen, printed, or saved as a text file. Figure A – 4 shows an example of a report of DHCP parameters, both global and per network segment.

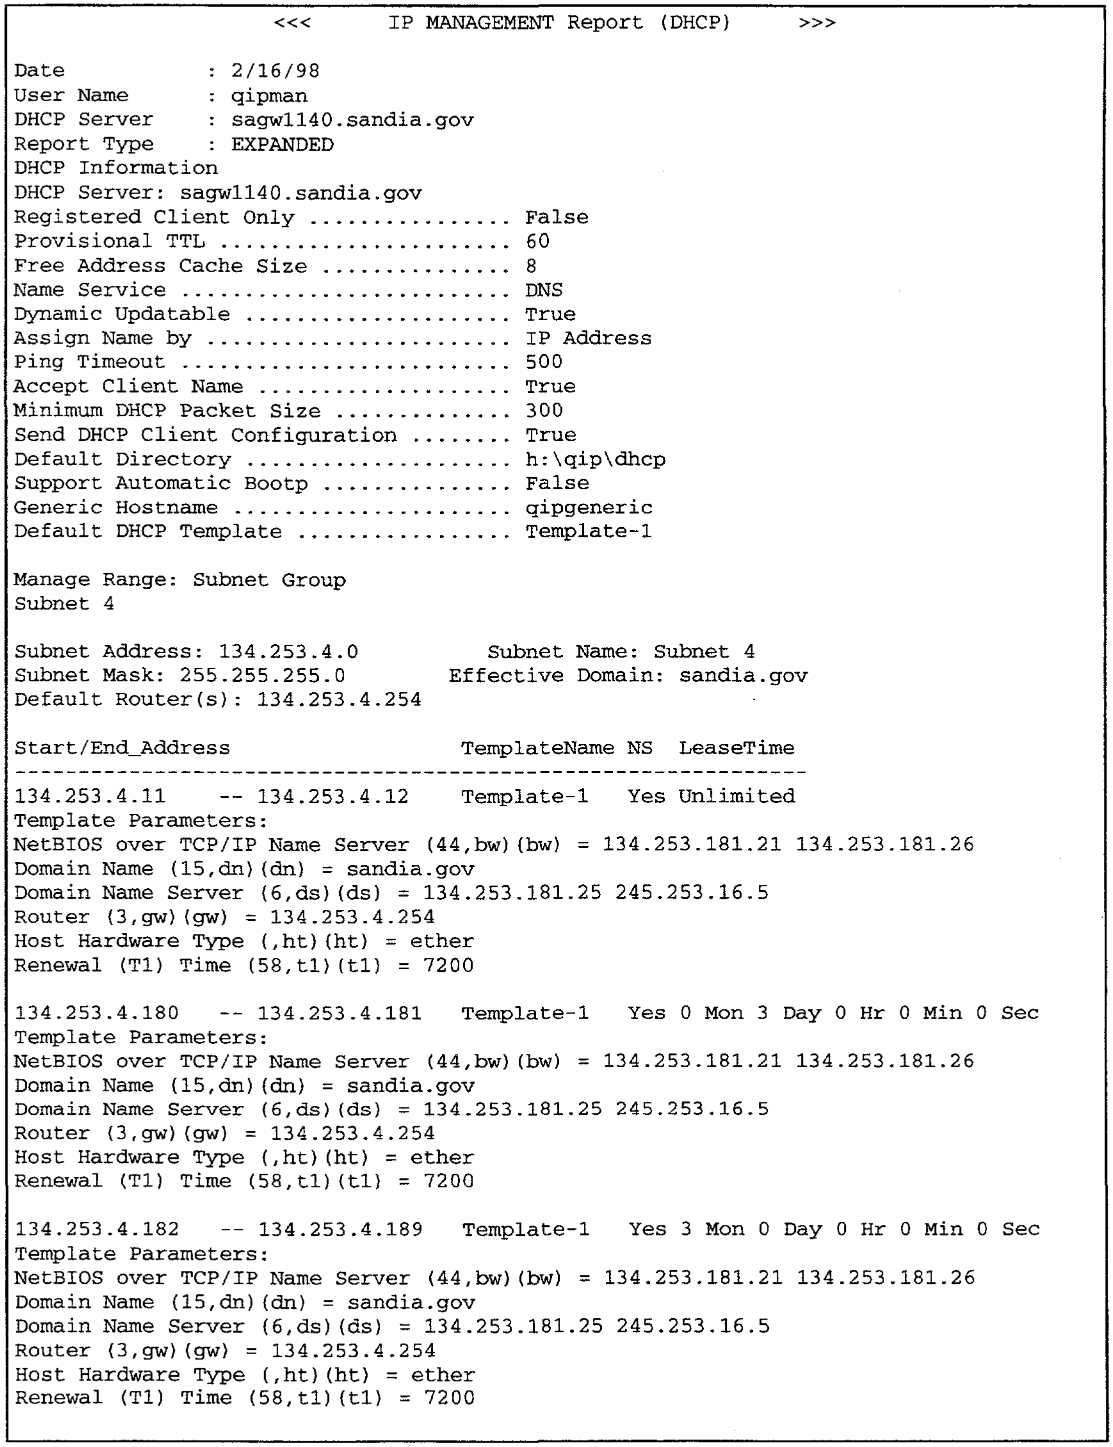

```
                    <<<      IP MANAGEMENT Report (DHCP)      >>>

Date           : 2/16/98
User Name      : qipman
DHCP Server    : sagw1140.sandia.gov
Report Type    : EXPANDED
DHCP Information
DHCP Server: sagw1140.sandia.gov
Registered Client Only ................ False
Provisional TTL ....................... 60
Free Address Cache Size ............... 8
Name Service .......................... DNS
Dynamic Updatable ..................... True
Assign Name by ........................ IP Address
Ping Timeout .......................... 500
Accept Client Name .................... True
Minimum DHCP Packet Size .............. 300
Send DHCP Client Configuration ........ True
Default Directory ..................... h:\qip\dhcp
Support Automatic Bootp ............... False
Generic Hostname ...................... qipgeneric
Default DHCP Template ................. Template-1

Manage Range: Subnet Group
Subnet 4

Subnet Address: 134.253.4.0         Subnet Name: Subnet 4
Subnet Mask: 255.255.255.0       Effective Domain: sandia.gov
Default Router(s): 134.253.4.254

Start/End_Address                  TemplateName NS  LeaseTime
-------------------------------------------------------------
134.253.4.11    -- 134.253.4.12    Template-1   Yes Unlimited
Template Parameters:
NetBIOS over TCP/IP Name Server (44,bw)(bw) = 134.253.181.21 134.253.181.26
Domain Name (15,dn)(dn) = sandia.gov
Domain Name Server (6,ds)(ds) = 134.253.181.25 245.253.16.5
Router (3,gw)(gw) = 134.253.4.254
Host Hardware Type (,ht)(ht) = ether
Renewal (T1) Time (58,t1)(t1) = 7200

134.253.4.180   -- 134.253.4.181   Template-1   Yes 0 Mon 3 Day 0 Hr 0 Min 0 Sec
Template Parameters:
NetBIOS over TCP/IP Name Server (44,bw)(bw) = 134.253.181.21 134.253.181.26
Domain Name (15,dn)(dn) = sandia.gov
Domain Name Server (6,ds)(ds) = 134.253.181.25 245.253.16.5
Router (3,gw)(gw) = 134.253.4.254
Host Hardware Type (,ht)(ht) = ether
Renewal (T1) Time (58,t1)(t1) = 7200

134.253.4.182   -- 134.253.4.189   Template-1   Yes 3 Mon 0 Day 0 Hr 0 Min 0 Sec
Template Parameters:
NetBIOS over TCP/IP Name Server (44,bw)(bw) = 134.253.181.21 134.253.181.26
Domain Name (15,dn)(dn) = sandia.gov
Domain Name Server (6,ds)(ds) = 134.253.181.25 245.253.16.5
Router (3,gw)(gw) = 134.253.4.254
Host Hardware Type (,ht)(ht) = ether
Renewal (T1) Time (58,t1)(t1) = 7200
```

Figure A - 4. IP Management Report (DHCP)

Event Logging

All activity generated by QIP is logged in the NT server event log. For example, when a client requests a lease (or a renewal), the event log will show the time and details of the request as shown in Figure A – 5.

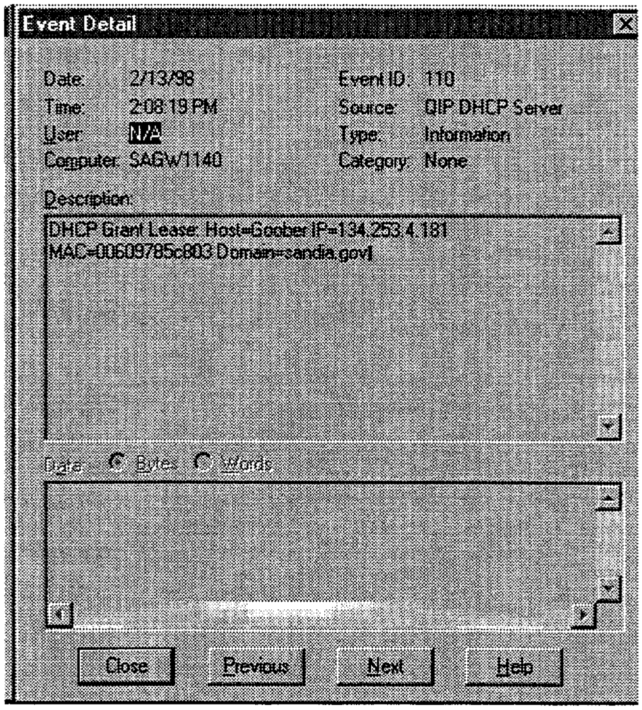

Date: 2/13/98
Time: 2:08:19 PM
User: N/A
Computer: SAGW1140

Event ID: 110
Source: QIP DHCP Server
Type: Information
Category: None

Description:

```
DHCP Grant Lease: Host=Goober IP=134.253.4.181
MAC=00609785c803 Domain=sandia.gov
```

Figure A - 5. Windows NT Event Log

Sample Active Leases File Dump

Displaying the active leases and choosing to create a text file generated the following file, Figure A - 6. This is the only product tested that could do this.

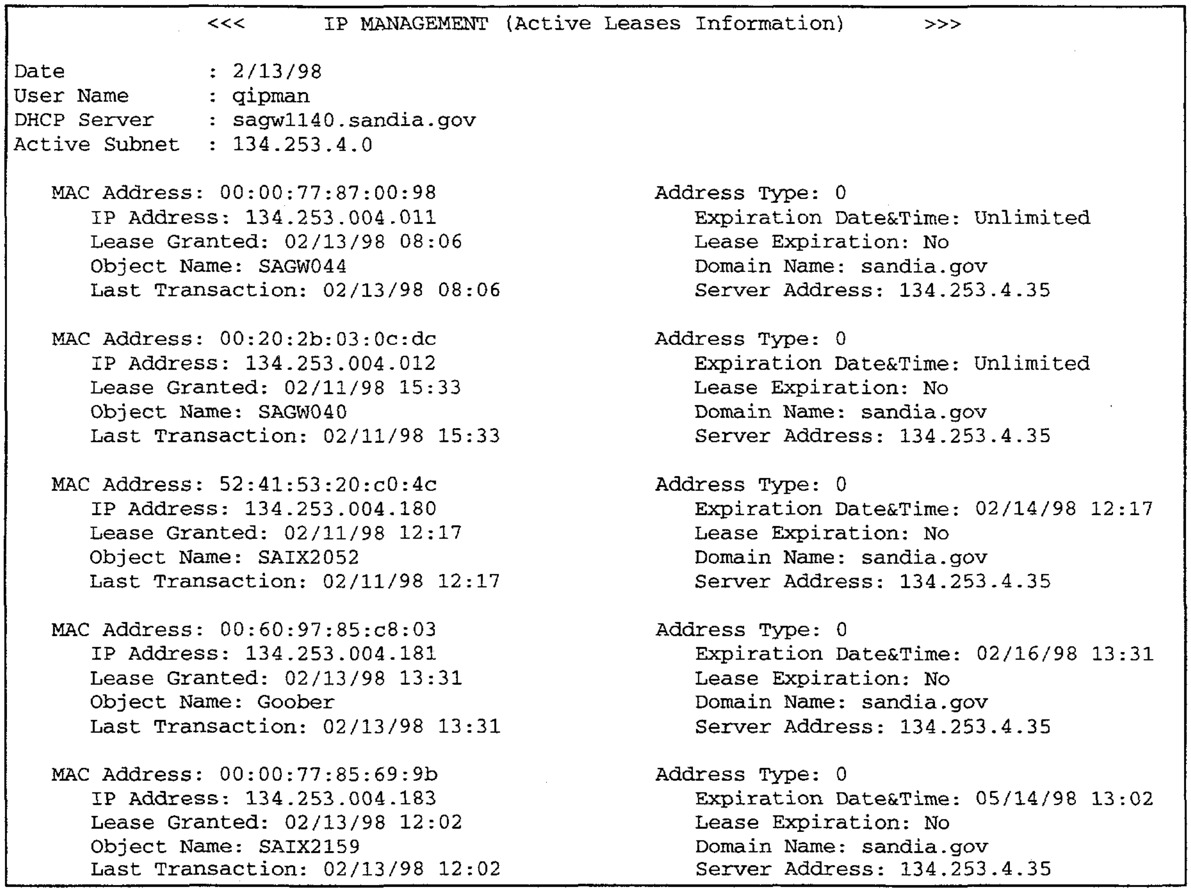

```
               <<<      IP MANAGEMENT (Active Leases Information)      >>>

Date          : 2/13/98
User Name     : qipman
DHCP Server   : sagw1140.sandia.gov
Active Subnet : 134.253.4.0

   MAC Address: 00:00:77:87:00:98                Address Type: 0
      IP Address: 134.253.004.011                   Expiration Date&Time: Unlimited
      Lease Granted: 02/13/98 08:06                 Lease Expiration: No
      Object Name: SAGW044                          Domain Name: sandia.gov
      Last Transaction: 02/13/98 08:06              Server Address: 134.253.4.35

   MAC Address: 00:20:2b:03:0c:dc                Address Type: 0
      IP Address: 134.253.004.012                   Expiration Date&Time: Unlimited
      Lease Granted: 02/11/98 15:33                 Lease Expiration: No
      Object Name: SAGW040                          Domain Name: sandia.gov
      Last Transaction: 02/11/98 15:33              Server Address: 134.253.4.35

   MAC Address: 52:41:53:20:c0:4c                Address Type: 0
      IP Address: 134.253.004.180                   Expiration Date&Time: 02/14/98 12:17
      Lease Granted: 02/11/98 12:17                 Lease Expiration: No
      Object Name: SAIX2052                         Domain Name: sandia.gov
      Last Transaction: 02/11/98 12:17              Server Address: 134.253.4.35

   MAC Address: 00:60:97:85:c8:03                Address Type: 0
      IP Address: 134.253.004.181                   Expiration Date&Time: 02/16/98 13:31
      Lease Granted: 02/13/98 13:31                 Lease Expiration: No
      Object Name: Goober                           Domain Name: sandia.gov
      Last Transaction: 02/13/98 13:31              Server Address: 134.253.4.35

   MAC Address: 00:00:77:85:69:9b                Address Type: 0
      IP Address: 134.253.004.183                   Expiration Date&Time: 05/14/98 13:02
      Lease Granted: 02/13/98 12:02                 Lease Expiration: No
      Object Name: SAIX2159                         Domain Name: sandia.gov
      Last Transaction: 02/13/98 12:02              Server Address: 134.253.4.35
```

Figure A - 6. IP Management (Active Leases Information)

Overall, this is a very nice product. It is a bit of a challenge to learn the configuration process. But once understood, this is no longer a problem. Client interaction seems to be adequate, except for the initial slow connect time for an old, slow laptop. No other client problems were encountered.

Sun Solaris DHCP

Features
DHCP server
Text-based database
No interaction with DNS
Remote administration by telnet session
Complex configuration

Over approximately a three week period, the Sun Solaris DHCP server software was tested. This is a bare-bones implementation. The address serving portion worked flawlessly, but the management portion is weak. The configuration software is a text-based program that prompts for parameters. The primary parameter is the range of IP addresses that are to be served up for the desired network segment. It takes only seconds to set up the server in its simplest form, but gets very difficult when considering other than default parameters (the usual UNIX documentation difficulty). For example, it is not clear if it is possible to establish exception addresses that are not to be served, or to configure ranges of addresses.

Logging of server activity is an important function that is given only cursory attention in this product. The status file for each configured network segment shows the addresses available for lease, the MAC address associated with all active leases, and the expiration date for each active lease. There is no information showing when the lease was last granted or when the address was last served. Knowing the lease time (typically three days by default), one can determine that an address was last issued sometime prior to the expiration time, if the expiration time is in the future. Otherwise, the address has not been served since the expiration time. This must all be inferred from the data.

Commands

- dhcpconfig - prompts administrator for parameters to configure dhcp
- dhtadm - DHCP configuration table management utility; allows administrator to modify or view the dhcp configuration table.
- pntadm - DHCP network table management utility; allows administrator to modify or view the dhcp network table.

This is a functional product, but not one that would suit our needs, due to the lack of administrative support.

Windows NT DHCP

Features
DHCP server only
No interaction with DNS
Proprietary remote administration
Simple Configuration
Proprietary database

This product was simple to configure and activate. It performed well (no performance problems were noted during testing). Remote administration is accomplished through a proprietary application that can be run from any Windows NT machine. A simple application appropriate for a small operation, but unsuitable for our needs due to lack of reporting, auditing, and dynamic DNS.

- Has a proprietary application for remote management that works well.

- No logging; current DHCP leases can be viewed on-line, but no record exists showing past lease history.
- Lease information is kept in a proprietary database (same type as a WINS database) which can not be viewed with typical tools such as Microsoft Access.
- The administrator can release leases.

Results

The conclusion reached from testing various DHCP products was that the QIP product best suited our needs. It met all mandatory requirements; server capable of serving multiple network segments, dynamic DNS, logging and auditing capability, and remote administration. It was also an easy product to use, which made it an attractive choice.

The Meta IP product came in as a close second. It is similar in function and appearance to QIP, but falls short in the logging and auditing department. For example, Meta IP can not generate a report of active leases, whereas QIP can not only generate a report, it can e-mail it, save it to disk, or print it. Another negative for Meta IP is that it uses Microsoft's Access for a database and this might not scale well in a large implementation.

The remaining products either did not meet mandatory requirements, or in the case of NetID, were too much trouble to evaluate in the very limited time available.

QIP it is!

Appendix B

Quadritek QIP Solution to Mobile Networking

Contents

Figures

APPENDIX B

Quadritek QIP Solution to Mobile Networking

The Quadritek QIP Enterprise DHCP/DNS product provides solutions to the problems of auditing and configuration control when providing network services to a nomadic worker within the campus environment. Of primary concern is the ability to determine who was responsible for particular network activity at some particular date and time. Addition concerns involve configuration control, such as who is able to make changes to configurations, and who made changes.

What computer had a given address at any time?

QIP Enterprise provides a display of active leases. These leases can be shown on screen for a quick look, saved to a text file, or printed.

The display of active leases would be the first place to start when attempting to determine what computer was using an address at some given time. Knowing the IP address, simply go to the Active Leases menu and choose the network segment that contains the address. Bring up the Active Leases Information screen and scan the list to see if the address in question appears. If it is in the list, then there will be a MAC address associated with the IP address.

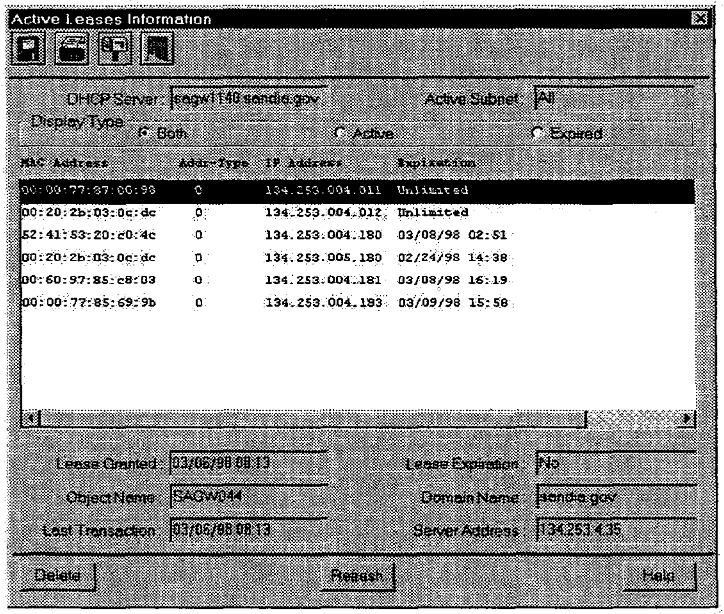

DHCP Server: sagw1140.sandia.gov
Active Subnet: All
Display Type: ● Both ○ Active ○ Expired

MAC Address	Addr-Type	IP Address	Expiration
00:00:77:87:00:98	0	134.253.004.011	Unlimited
00:20:2b:03:0c:dc	0	134.253.004.012	Unlimited
52:41:53:20:c0:4c	0	134.253.004.180	03/08/98 02:51
00:20:2b:03:0c:dc	0	134.253.005.180	02/24/98 14:38
00:60:97:85:c8:03	0	134.253.004.181	03/08/98 16:19
00:00:77:85:69:9b	0	134.253.004.183	03/09/98 15:58

Lease Granted: 03/06/98 08:13
Lease Expiration: No
Object Name: SAGW044
Domain Name: sandia.gov
Last Transaction: 03/06/98 08:13
Server Address: 134.253.4.35

Figure B - 1. Active Leases Information Screen

Figure B – 1 is a screen shot of active leases (as well as expired ones that have not been reassigned) for a particular set of network segments (134.253.4 and 134.253.5). Radio buttons provide the ability to show expired leases or filter them out to show only the active leases. For a large network segment, a vertical scroll bar appears to allow scrolling through the list of leases, sorted by IP address. Details of the selected row appear at the bottom of the screen. In this example, the active lease for 134.253.4.11 is associated with SAGW044.SANDIA.GOV, with the last DHCP transaction occurring on 3/6/98. The computer name, as shown in the Object Name field, is obtained from the DHCP client computer during the DHCP lease grant or renewal.

If this IP address was in question on 3/7/98, looking at the Active Leases Information screen would show that this address was associated with MAC address 00:00:77:87:00:98, and we can then attempt to determine the owner by looking up this MAC address in Sandia's Network Information System (NWIS).

The text version of the display of active leases is shown in Figure B - 2 and is suitable for including in reports or for printing. It includes all the same information as the screen version.

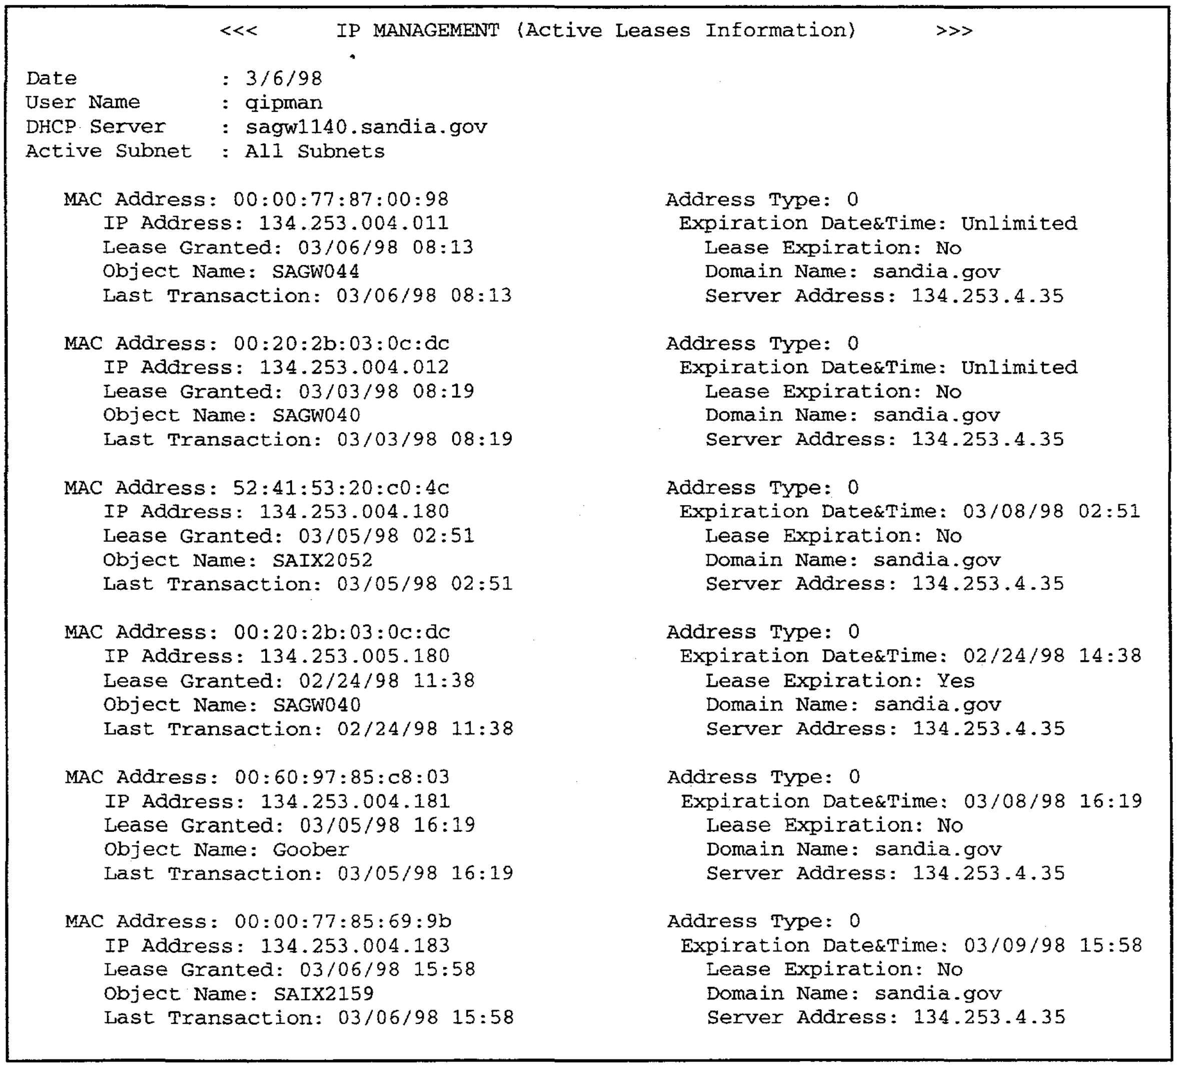

```
                <<<      IP MANAGEMENT (Active Leases Information)      >>>

Date           : 3/6/98
User Name      : qipman
DHCP Server    : sagw1140.sandia.gov
Active Subnet  : All Subnets

   MAC Address: 00:00:77:87:00:98                  Address Type: 0
      IP Address: 134.253.004.011                   Expiration Date&Time: Unlimited
      Lease Granted: 03/06/98 08:13                   Lease Expiration: No
      Object Name: SAGW044                            Domain Name: sandia.gov
      Last Transaction: 03/06/98 08:13                Server Address: 134.253.4.35

   MAC Address: 00:20:2b:03:0c:dc                  Address Type: 0
      IP Address: 134.253.004.012                   Expiration Date&Time: Unlimited
      Lease Granted: 03/03/98 08:19                   Lease Expiration: No
      Object Name: SAGW040                            Domain Name: sandia.gov
      Last Transaction: 03/03/98 08:19                Server Address: 134.253.4.35

   MAC Address: 52:41:53:20:c0:4c                  Address Type: 0
      IP Address: 134.253.004.180                   Expiration Date&Time: 03/08/98 02:51
      Lease Granted: 03/05/98 02:51                   Lease Expiration: No
      Object Name: SAIX2052                           Domain Name: sandia.gov
      Last Transaction: 03/05/98 02:51                Server Address: 134.253.4.35

   MAC Address: 00:20:2b:03:0c:dc                  Address Type: 0
      IP Address: 134.253.005.180                   Expiration Date&Time: 02/24/98 14:38
      Lease Granted: 02/24/98 11:38                   Lease Expiration: Yes
      Object Name: SAGW040                            Domain Name: sandia.gov
      Last Transaction: 02/24/98 11:38                Server Address: 134.253.4.35

   MAC Address: 00:60:97:85:c8:03                  Address Type: 0
      IP Address: 134.253.004.181                   Expiration Date&Time: 03/08/98 16:19
      Lease Granted: 03/05/98 16:19                   Lease Expiration: No
      Object Name: Goober                             Domain Name: sandia.gov
      Last Transaction: 03/05/98 16:19                Server Address: 134.253.4.35

   MAC Address: 00:00:77:85:69:9b                  Address Type: 0
      IP Address: 134.253.004.183                   Expiration Date&Time: 03/09/98 15:58
      Lease Granted: 03/06/98 15:58                   Lease Expiration: No
      Object Name: SAIX2159                           Domain Name: sandia.gov
      Last Transaction: 03/06/98 15:58                Server Address: 134.253.4.35
```

Figure B - 2. Active Leases Information Report

If the address in question is not found in the Active Leases Information screen, or is determined to be from a different computer than what is shown as active, then it will be necessary to scan the NT Server event log (and/or its archives). All QIP DHCP activity is entered into the event log.

Figure B - 3 is an example of an NT server event log, filtered to show only DHCP lease activity. In this form, it is very difficult to search for specific data within an event. The viewer shows the events, but has no search capability for finding specific information such as when a particular IP address was involved in a lease grant or renewal. Individual events must be opened in order to see the details of the event.

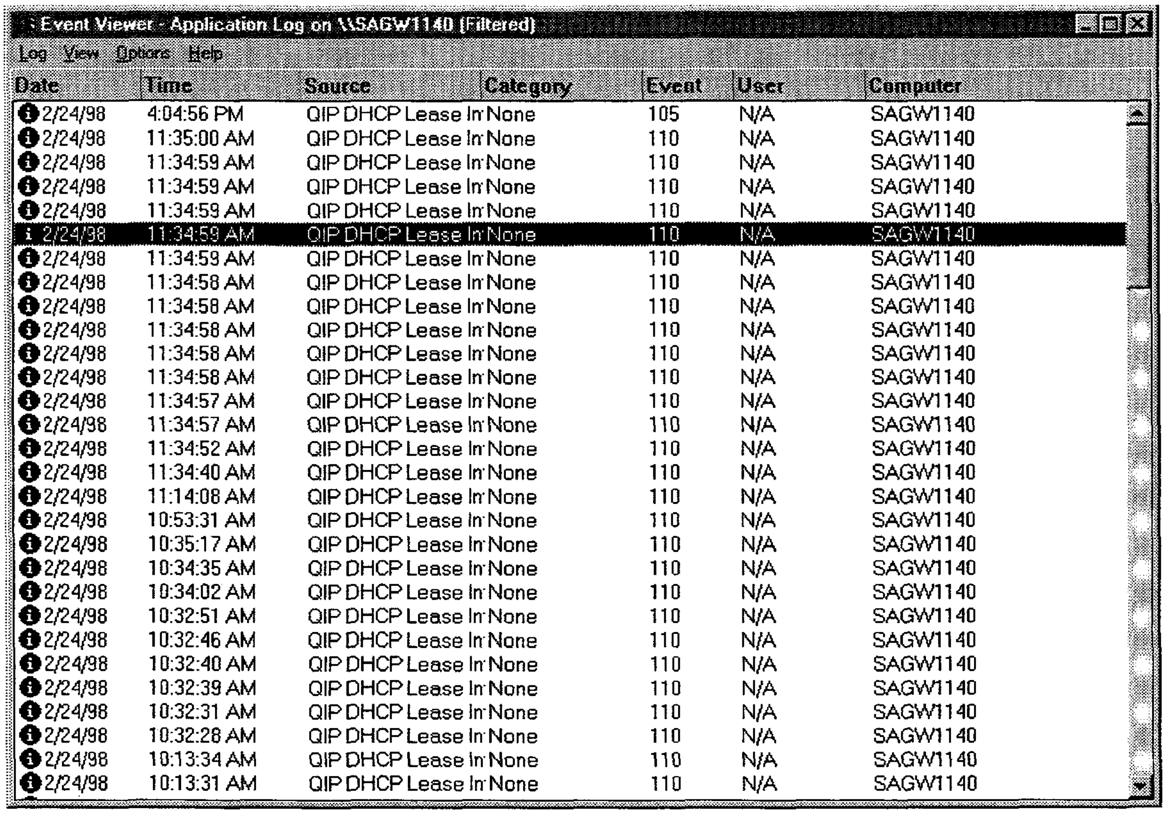

Date	Time	Source	Category	Event	User	Computer
ℹ 2/24/98	4:04:56 PM	QIP DHCP Lease In	None	105	N/A	SAGW1140
ℹ 2/24/98	11:35:00 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:59 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:59 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:59 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:59 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:59 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:58 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:58 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:58 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:58 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:58 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:57 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:57 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:52 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:34:40 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	11:14:08 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:53:31 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:35:17 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:34:35 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:34:02 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:32:51 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:32:46 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:32:40 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:32:39 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:32:31 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:32:28 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:13:34 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140
ℹ 2/24/98	10:13:31 AM	QIP DHCP Lease In	None	110	N/A	SAGW1140

Figure B - 3. Windows NT Event Log Showing Lease Activity

How to deal with the event log?

Searching the NT event log for DHCP lease information can be facilitated by importing the event log into a spreadsheet. This is actually a two step process where the event log is first exported to a comma delimited text file, then imported into the spreadsheet. Once in the spreadsheet, the data can be searched or sorted, and reports can be generated.

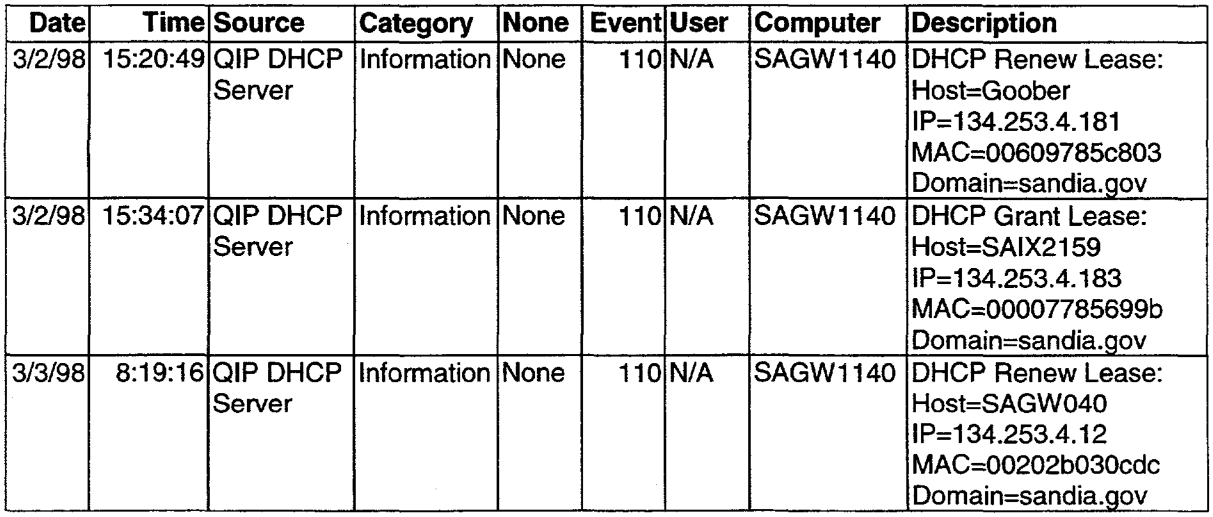

Date	Time	Source	Category	None	Event	User	Computer	Description
3/2/98	15:20:49	QIP DHCP Server	Information	None	110	N/A	SAGW1140	DHCP Renew Lease: Host=Goober IP=134.253.4.181 MAC=00609785c803 Domain=sandia.gov
3/2/98	15:34:07	QIP DHCP Server	Information	None	110	N/A	SAGW1140	DHCP Grant Lease: Host=SAIX2159 IP=134.253.4.183 MAC=00007785699b Domain=sandia.gov
3/3/98	8:19:16	QIP DHCP Server	Information	None	110	N/A	SAGW1140	DHCP Renew Lease: Host=SAGW040 IP=134.253.4.12 MAC=00202b030cdc Domain=sandia.gov

Figure B - 4. Windows NT Event Log Data in Spreadsheet

Figure B – 4 shows an example of a spreadsheet generated from an event log. Note that only DHCP server messages appear in the spreadsheet. Although the NT event log stores many types of events, both DHCP and other types, a filter was created within the spreadsheet to show only DHCP server messages. Now a particular date and time can be located with respect to DHCP lease activity. Or A particular IP address can be located to show lease activity.

Who made changes to DHCP configurations?

Figure B - 5 shows the history for the primary administrator of the QIP product. The data can be saved as a file, printed, or even e-mailed. The screen shows the action taken, the object that was acted on, the IP address, object status, object class, and the date of the action. Note that in the case of a command to perform a modification, the old and new data is shown. For example, on this screen, object SAGW040 was changed from old object name ws0069 to new object name SAGW040 on 2/17/89 at 2:27 p.m.

If other administrators are defined in the system, the administrator name can be input and the history report generated.

APPENDIX B

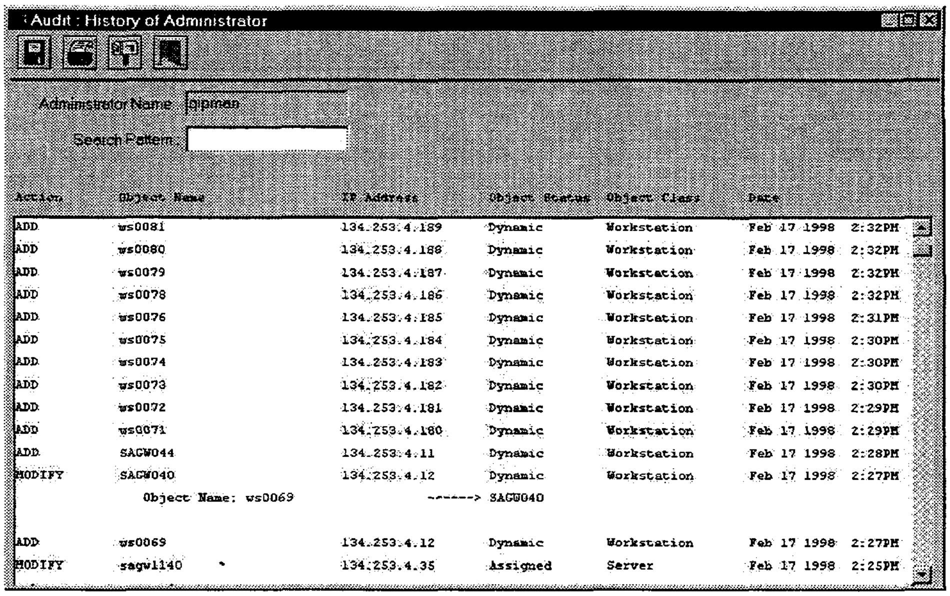

Audit : History of Administrator

Administrator Name: qipman

Search Pattern:

Action	Object Name	IP Address	Object Status	Object Class	Date
ADD	ws0081	134.253.4.189	Dynamic	Workstation	Feb 17 1998 2:32PM
ADD	ws0080	134.253.4.188	Dynamic	Workstation	Feb 17 1998 2:32PM
ADD	ws0079	134.253.4.187	Dynamic	Workstation	Feb 17 1998 2:32PM
ADD	ws0078	134.253.4.186	Dynamic	Workstation	Feb 17 1998 2:32PM
ADD	ws0076	134.253.4.185	Dynamic	Workstation	Feb 17 1998 2:31PM
ADD	ws0075	134.253.4.184	Dynamic	Workstation	Feb 17 1998 2:30PM
ADD	ws0074	134.253.4.183	Dynamic	Workstation	Feb 17 1998 2:30PM
ADD	ws0073	134.253.4.182	Dynamic	Workstation	Feb 17 1998 2:30PM
ADD	ws0072	134.253.4.181	Dynamic	Workstation	Feb 17 1998 2:29PM
ADD	ws0071	134.253.4.180	Dynamic	Workstation	Feb 17 1998 2:29PM
ADD	SAGW044	134.253.4.11	Dynamic	Workstation	Feb 17 1998 2:28PM
MODIFY	SAGW040	134.253.4.12	Dynamic	Workstation	Feb 17 1998 2:27PM
	Object Name: ws0069	------>	SAGW040		
ADD	ws0069	134.253.4.12	Dynamic	Workstation	Feb 17 1998 2:27PM
MODIFY	sagw1140	134.253.4.35	Assigned	Server	Feb 17 1998 2:25PM

Figure B - 5. Audit Screen - History of Administrator

Figure B - 6 an example of a history report on an object, rather than an administrator. This can be very useful for narrowing down activity of an object, rather than having to sort through administrators to find data on a specific object.

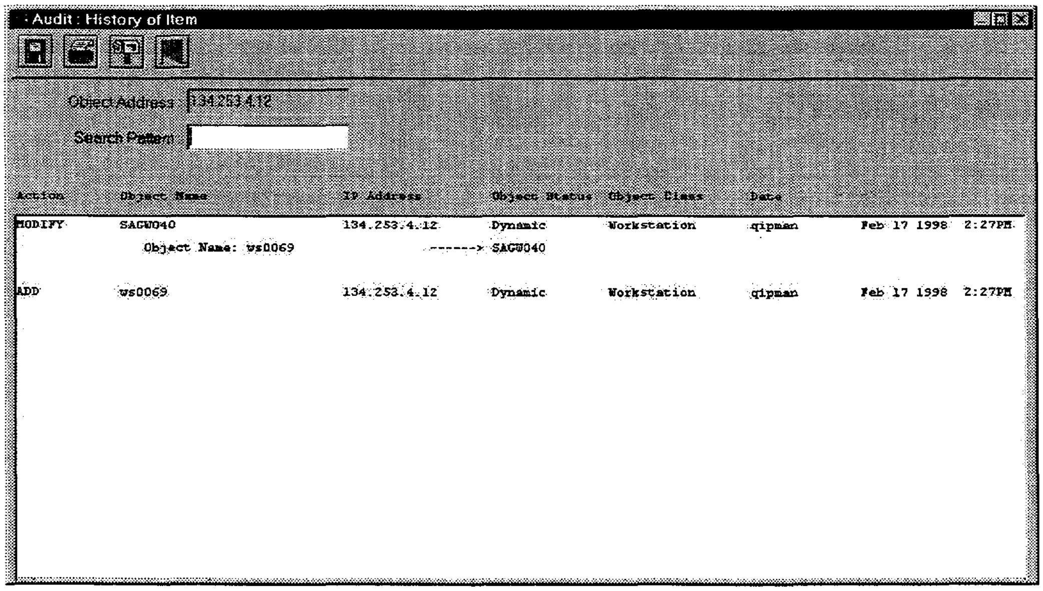

Audit : History of Item

Object Address: 134.253.4.12

Search Pattern:

Action	Object Name	IP Address	Object Status	Object Class	Date	
MODIFY	SAGW040	134.253.4.12	Dynamic	Workstation	qipman	Feb 17 1998 2:27PM
	Object Name: ws0069	------>	SAGW040			
ADD	ws0069	134.253.4.12	Dynamic	Workstation	qipman	Feb 17 1998 2:27PM

Figure B - 6. Audit Screen - History of Item

Who can make changes to DHCP configurations?

QIP Enterprise has a hierarchical system for defining administrators. The primary administrator must be defined and has authority over all aspects of the data. Secondary administrators can be defined that have authority only over their subset of modules, such as an administrator that has responsibility over a particular network segment and no others.

Figure B - 7 shows the primary administrator definitions. The account can be tailored to provide a default printer and warning conditions.

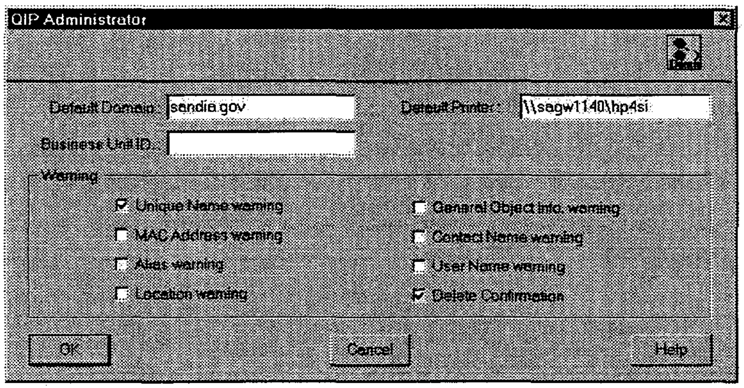
QIP Administrator

Default Domain: sandia.gov
Default Printer: \\sagw1140\hp4si
Business Unit ID:

Warning
☑ Unique Name warning
☐ MAC Address warning
☐ Alias warning
☐ Location warning
☐ General Object info. warning
☐ Contact Name warning
☐ User Name warning
☑ Delete Confirmation

Figure B - 7. Primary Administrator Definitions

Figure B - 8 shows a network segment administrator defined for the network segment 134.253.4.0. This administrator can only administer this network segment and the addresses specified in the Address Range (Start) and (Stop) fields. The granularity of this product allows administrators to be assigned to different ranges of addresses within the same network segment; for example, a network segment might be divided between different organizations.

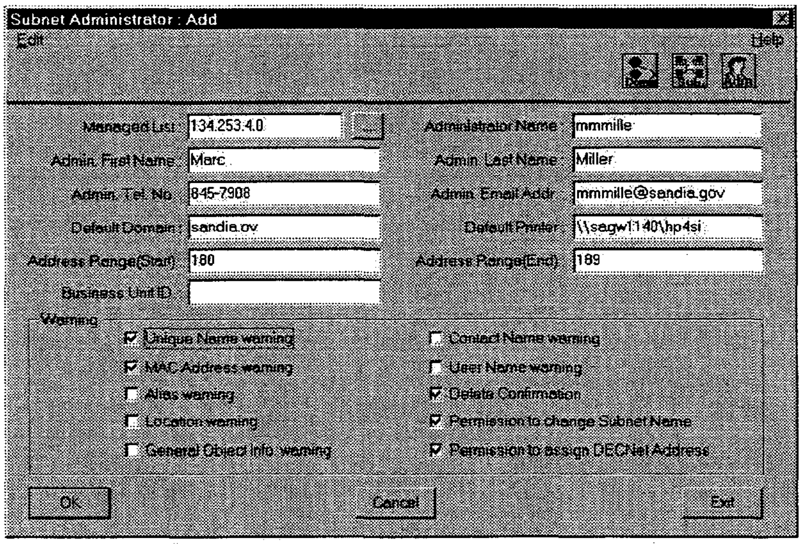
Subnet Administrator : Add

Managed List: 134.253.4.0
Administrator Name: mmmille
Admin. First Name: Marc
Admin. Last Name: Miller
Admin. Tel. No: 845-7908
Admin. Email Addr: mmmille@sandia.gov
Default Domain: sandia.ov
Default Printer: \\sagw1140\hp4si
Address Range(Start): 180
Address Range(End): 189
Business Unit ID:

Warning
☑ Unique Name warning
☑ MAC Address warning
☐ Alias warning
☐ Location warning
☐ General Object Info. warning
☐ Contact Name warning
☐ User Name warning
☑ Delete Confirmation
☑ Permission to change Subnet Name
☑ Permission to assign DECNet Address

Figure B - 8. Subnet Administrator Screen

Integrated Event Analysis Available with Next Release of QIP

The previously described procedures for auditing and configuration control are certainly not attractive, but nonetheless necessary. They actually add administrative burden, which is counter to the desire to reduce these types of tasks. Quadritek recognizes this problem and has addressed it with the next release, QIP Enterprise 5.0. This release will include the Audit Manager, an analysis and planning module that will provide network and security administrators with centralized tools to track address activities among all DHCP and DNS servers.

Additional modules analyze and manage data related to the state and use of multiple IP services. The Policy Manager allows network planners to establish and enforce rules consistently throughout the network for how IP services are to function and inter-operate. The Profile Manager provides the design tools to configure the network, and define the relationships among network objects. The Services Manager allows network managers to view, analyze, and interact with the IP services on remote servers throughout the network.

These tools, coupled with automated IP address management using DHCP, should create an efficient system for administrating a large IP network.

Appendix C

DHCP Client Configuration

DHCP Client Configuration

Windows NT

To configure a Windows NT machine to operate with DHCP, bring up the Network Properties dialog box by right clicking on the Network Neighborhood desktop icon, or by selecting Network in the Control Panel. Go to the Protocols tab, select the TCP/IP Protocol and either double click or select the Properties button. In the Microsoft TCP/IP Properties dialog box, as shown in Figure C - 1, select the IP Address tab. Select the button for "Obtain an IP address from a DHCP server". Click the OK button and reboot as prompted. This is the only change necessary to enable a computer to utilize DHCP. However, there are many other issues involving the client configuration for Windows NT.

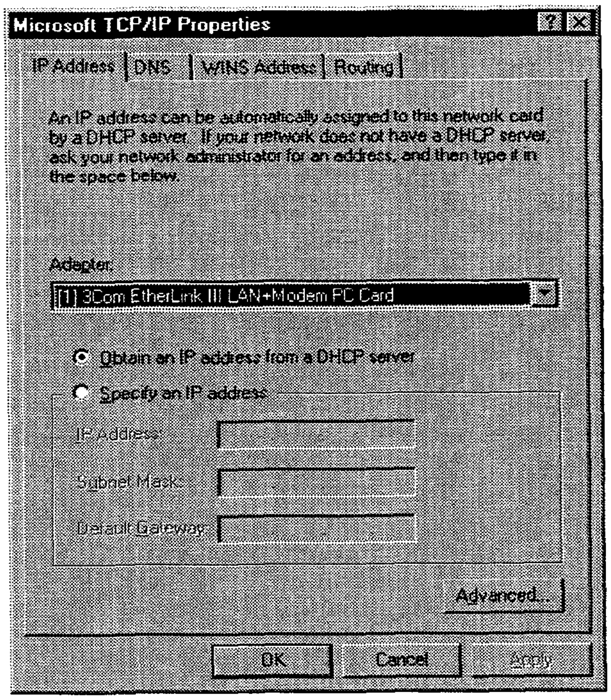

An IP address can be automatically assigned to this network card by a DHCP server. If your network does not have a DHCP server, ask your network administrator for an address, and then type it in the space below.

Adapter:
[1] 3Com EtherLink III LAN+Modem PC Card

● Obtain an IP address from a DHCP server
○ Specify an IP address

IP Address:
Subnet Mask:
Default Gateway:

Figure C - 1. Windows NT TCP/IP Properties

DNS, WINS, and router (or gateway) parameters may or may not be set in the client. None of these parameters is necessary for DHCP operation, assuming the DHCP server is configured to supply the parameters. However, they can have an effect. If the parameters are left blank, the DHCP server will provide them to the client. If the parameters are set, they will override the values from the DHCP server. This is probably of little significance except in cases where a client is using a local DNS server, then configures to use DHCP and removes the local DNS entry, and subsequently starts using the DNS server supplied by DHCP. This should not happen often as most client machines us the corporate DNS servers. The same goes for the WINS servers.

On the DNS configuration screen, there is a place to enter the machine name, called "Host Name." This box is normally filled in and any value entered here will override the value supplied by DHCP. Again, however, this field could be left blank, especially in the case of a new machine that is configured to use DHCP and no other parameters are entered. If left blank, the DHCP server will give the machine a name. This name will be some random-looking string that is guaranteed to be unique within the DHCP realm, but will be totally meaningless to the owner of the machine. Therefore, it is always a good idea to give the machine a name (preferably, the name assigned by NWIS, the Sandia Network Information System).

Also on the DNS configuration screen is a field for the domain. This is the DNS domain (not to be confused with a Windows network domain). This field, like the name field, can be left blank and will be assigned by the DHCP server. In fact, it is better to leave the Domain field blank and let DHCP do its job. However, it if this field is filled in, most likely with SANDIA.GOV, then there will be no harm done.

To set the gateway for Windows NT, select the Advanced button on the TCP/IP Properties screen. The Gateways box can contain one or more router addresses and may or may not have any entries.

Windows 95

As with Windows NT, the Windows 95 TCP/IP Properties dialog box, shown in Figure C - 2, is accessed by either right clicking on the Network Neighborhood desktop icon, or by double clicking the Network icon in the control panel. Select the network adapter from the list and either double click or select the Properties button. To enable DHCP, simply select the button "Obtain an IP address automatically" and reboot. As with NT, the DNS, WINS, and router configuration screens are optional and have the same effect as with NT.

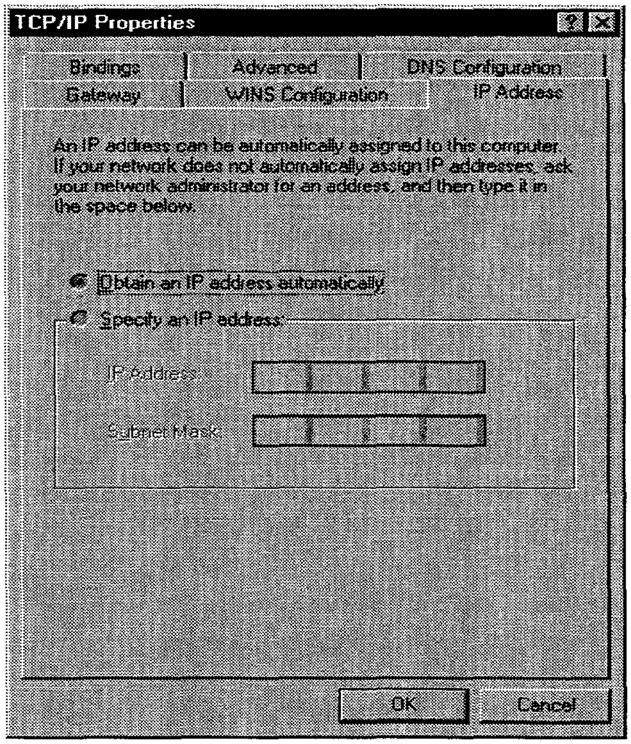

An IP address can be automatically assigned to this computer. If your network does not automatically assign IP addresses, ask your network administrator for an address, and then type it in the space below.

● Obtain an IP address automatically

○ Specify an IP address:

Figure C - 2. Windows 95 TCP/IP Properties

Distribution

1	MS	0449	R. S. Tamashiro, 6237
1		0469	J. F. Jones, 4600
1		0622	D. L. Weaver, 4013
1		0629	P. D. Merillat, 4800
1		0630	J. P. Vandevender, 4010
1		0662	L. B. Cox, 4421
1		0662	T. Klitsner, 4423
1		0662	D. S. Rarick, 4423
1		0801	M. J. Murphy, 4400
1		0803	L. H. Pitts, 4900
1		0805	W. F. Chambers, 4911
1		0805	R. N. Harris, 4911
1		0806	C. D. Brown, 4621
1		0806	D. E. Ellis, 4621
1		0806	S. A. Gossage, 4616
1		0806	J. A. Hudson, 4616
1		0806	J. P. Long, 4621
1		0806	G. D. Machin, 4621
15		0806	M. M. Miller, 4616
1		0806	P. C. Moore, 4621
1		0806	L. G. Pierson, 4616
1		0806	M. O. Vahle, 4616
1		0806	E. L. Witzke, 4616
1		0807	I. C. Alexander, 4417
1		0807	S. D. Nelson, 4417
1		0811	T. L. Ferguson, 4813
1		0812	M. R. Sjulin, 4914
1		0812	L. F. Tolendino, 4914
1		0812	B. C. Whittet, 4914
1		0817	L. D. Daigle, 4815
1		1202	D. A. Rieb, 5931
1		9011	P. R. Bryson, 8910
1		9011	H. Y. Chen, 8910
1		9011	P. W. Dean, 8903
1		9011	D. A. Evensky, 8980
1		9012	F. T. Bielecki, 8930
1		9012	R. D. Gay, 8930
1		9012	S. C. Gray, 8930
1		9018	Central Technical Files, 8940-2
2		0899	Technical Library, 4916
1		0619	Review & Approval Desk, 15102 For DOE/OSTI